

PATENT APPLICATION
DOCKET NO. 0100.9901410

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

FILING OF A UNITED STATES PATENT APPLICATION

TITLE:

METHOD AND SYSTEM FOR HANDLING ERRORS

INVENTORS:

**Branko Kovacevic
Kevork Kechichian**

ATTORNEY OF RECORD

**J. GUSTAV LARSON
REGISTRATION NO. 39,263**

**P.O. Box 06229
CHICAGO, ILLINOIS 60606-00229
PHONE (512) 336-8957
FAX (512) 336-9155**

EXPRESS MAIL LABEL NO. **EL394829446US**

DATE OF DEPOSIT: **1/22/00 1/24/00 pm**

I HEREBY CERTIFY THAT THIS PAPER IS BEING
DEPOSITED WITH THE U.S. POSTAL SERVICE "EXPRESS
MAIL POST OFFICE TO ADDRESSES" SERVICE UNDER 37
C.F.R. SECTION 1.10 ON THE 'DATE OF DEPOSIT',
INDICATED ABOVE, AND IS ADDRESSED TO THE
COMMISSIONER OF PATENTS AND TRADEMARKS,
WASHINGTON, D.C. 20231.

NAME OF DEPOSITOR:

(PRINT OR TYPE)

SIGNATURE:

Jane Markov

00485669-012400

METHOD AND SYSTEM FOR HANDLING ERRORS

Field of the Invention

The present invention relates generally to detecting and handling errors in the packetized data, and more specifically to detection and handling of MPEG-2 transport stream errors.

Background of the invention

The international organization for standards (ISO) moving pictures experts group (MPEG group), approved an audio video digital compression standard known as MPEG-2 in an effort to provide a versatile compression standard capable of being utilized for a wide variety of data. The MPEG-2 standard provides explanations needed to implement an MPEG-2 decoder through the use of syntax and semantics of a coded bit stream. MPEG-2 is an open standard which continues to evolve and be applied to a variety of applications ranging from video conferencing to high definition television. As a generic standard, MPEG-2 is intended to be used for variety of audio and video coding applications. Part one of the MPEG-2 standard (ISO 13818-1), was designated to improve error resilience and carry multiple programs simultaneously without a common time base between programs.

1221

The transport stream (TS) specified by the MPEG-2 standard, offers a high degree of robustness for noisy channels, and can be used to carry multiple programs, such as multiple TV services. The transport stream is based on a 188 byte long packet suited for hardware error correction and processing schemes. The use of a robust protocol, such as the transport stream, allows for reception over noisy environments such as terrestrial and satellite transmissions. Even in these environments it is possible to obtain fast program access, channel hoping, and synchronization between multiple elementary streams carried within the packetized elementary streams which are subdivided into transport packets.

Prior art Figure 1 illustrates a Transport Packet Stream defined by the MPEG-2 standard. The transport stream, based on a 188 byte long packet, is well suited for hardware error correction and processing schemes. Such a configuration can carry multiple programs within the same multiplex, even when the transmission environment is noisy. For example, MPEG-2 data can be transferred successfully over coaxial cable networks and satellite transponders with asynchronous multiplexing of constant or variable bit-rate programs to allow fast program access, channel hopping and synchronization between services.

As illustrated further in Figure 1, MPEG-2 transport stream consists of fixed length Transport Stream Packets (TSP or packets) based on 4 bytes of header followed by 184 bytes of TSP payload. TSP payload carries Packetized Elementary Stream (PES) data obtained by chopping up an Elementary Stream (ES), which consists of data of a common type and program. For example, audio for a specific program would form one elementary stream, while video for the same program would form a second elementary stream.

The TS header consists of a synchronization byte (SyncByte), flags, information indicators for error detection and timing, an adaptation field indicator, and a Packet_ID (PID) field used to identify Elementary Streams carried in the payload. The adaptation field, when present, contains flags, and timing information.

The PID Field is used not only to distinguish separate Elementary Streams, but also separate Program Specific Information (PSI) tables. Prior art Figure 2 illustrates two types of PSI tables -- a Program Association Table 210 (PAT), and a Program Map Table 220 (PMT). The PAT table lists unique program numbers as identifiers for each program, or elementary stream, in a multiplex, and the PID number associated with each program number. A fixed PID number of 0x0000 is assigned to the PAT table, making it possible for the system to download the PAT table on startup by retrieving PID 0x0000 packets.

Each program identified the PAT table has a related Program Map Table (PMT) having its own PID identifier. Each PMT table lists the PIDs for all Elementary Streams (components) making a given program associated with the PMT. A specific PMT table

maybe constructed for each program separately, or may be common for a group of programs. In the first case, there are many PMT tables with just one section, and each PMT table has a different PID value. In the second case one PMT table may have many sections, each relevant to one program.

In order to provide multiple services over the same multiplex, data associated with different multimedia services are transmitted, with packet multiplexing, such that data packets from several Elementary Streams of audio, video, data, and others are interleaved on a packet by packet basis into a single MPEG-2 transport stream. Synchronization between Elementary Streams forming a common program is achieved using presentation time stamps and program clock references which can be transmitted as part of the adaptation field specified in the header.

Prior art Figure 3 illustrates the fields associated with a PES stream. Each PES stream contains a header portion and a data portion. In addition, an optional header portion may exist. The header portion includes a Packet Start Prefix, a stream ID, and a packet length indicator.

Transport stream information can be provided either through a direct broadcast, or through a service provider broadcast. Direct broadcast generally refers to signals which are received directly by an end user. Examples of direct broadcasts include satellite broadcasts received by satellite dishes and provided to a decoder at the end user's location, which receives and decodes the transport stream data. Another type of direct broadcast is the traditional composite television/radio broadcast. In their most elementary forms, these broadcasts are not digital broadcasts. However, the transmission of digital broadcast in MPEG-2 format is being explored and occurring as an alternative. In this manner, the user would have a tuner capable of receiving the direct terrestrial link information containing the television or radio signals. Once demodulated, the transport stream information could be provided to a desktop unit, or decoder, owned by the end user.

Service provider broadcast would include broadcast to the home provided by cable television providers, telephone company providers, or other independent providers.

In this configuration, the service provider first receives the number of signals which can be ultimately provided to the end user. Examples of such received signals include satellite feeds, terrestrial feeds, switched video sources, local video sources such as tapes, or laser disk DVD's, as well as traditional table feeds. Based upon the end users demands, the received information can be selectively provided to the end user.

In one manner, the selected feed by the service provider can be provided directly to an end user through a twisted pair connection, which may include a high speed digital subscriber link (DSL) capable of providing data at higher rates than traditionally associated with plain old telephone system (POTS) connections.

In another implementation, the service provider would provide information from a central office or a head-end to a fiber node. A specific fiber node is generally used to support more than one end user. Examples of the use of such fiber nodes includes a fiber coaxial bus (FCB) whereby a fiber link provides the signal containing a large amount of data to a fiber node which in turn drives coaxial cable having a taps. A decoding device attached to taps at user side can receive the appropriate broadcasting signal.

Another example of a fiber node is bi-directional fiber coaxial bus. While similar to the FCB bus, the bi-directional FCB bus is also capable of transmitting data back to the central office or the head-end as well as receiving it. Yet another fiber node example is a hybrid fiber coax, which uses coaxial cable and branch topology toward network interface units. Yet another example associated with service providers is known as fiber to the curb, whereby digital signaling is carried from the central office to the optical network unit which serves only a few dozen homes. Locally, a hybrid twisted pair coaxial pairs will connect to the optical network unit with the consumer's decoder. Twist repair cable carries digital video in the 5 to 40 megahertz range to no more than 500 feet from the fiber connection. Therefore, the number of homes capable of being served by a single optical network unit is limited. Analog television signals are carried in a coaxial cable from the fiber node.

One problem associated with the flexibility of the MPEG-2 standard is that once the transport stream is received, demodulated, and decrypted, the resulting data stream

g2
amt

can still have a variety variations which need be known before the data stream can be properly utilized. For example, the MPEG-2 specification does not indicate a specific set of control signals to be provided with the transport stream, how received data and control signals are qualified, nor the precise format of the actual data transmitted. As a result, implementations of set top boxes require specific service provider information. Specific service provider information results in an incompatibility among transport streams schemes provided by different service providers or cable operators. As a result, chip sets are designed and dedicated to support specific service provider's set top boxes.

Prior art Figure 4 illustrates a prior art system for parsing a transport stream. The transport parser of the prior art would receive individual packets from the framer. Based upon the PID value, the transport parser would store the TSP data to be used by the system or the graphics engine in a local buffer.

When the transport parser's local buffer was filled, the transport parser would cause a bus request to the appropriate controller (system or video) to initialize a transfer of at least some of the buffered data.

However, when the prior art host video or graphics system needed more data from the transport parser prior to the transport parser initializing the transfer, the system would initialize the transfer by generating a request to access data in the transport parser buffer. Since the bus used internally by the transport parser buffer may have other clients, the host system may have to wait to access the bus. The overall performance of the host system is reduced as a result of the system waiting on data.

Therefore, a system and method of receiving transport stream information that allows for more flexibility and improved performance in terms of data handling, data parsing, design implementation, data stream acquisition would be advantageous.

Brief Description of the Drawings

Figure 1 illustrates, in block form, prior art fields associated with a transport stream packet;

Figure 2 illustrates, in tabular form, a prior art Program Specific Information tables;

Figure 3 illustrates, in block form, prior art fields associated with Packetized Elementary Stream;

Figure 4 illustrates, a prior art representation of a parser system;

Figure 5 illustrates, in block diagram form, a transport stream core in accordance with the present invention;

Figure 6 illustrates a tabular representation of a register set;

Figure 7 illustrates, in block diagram form, another embodiment of a transport stream core in accordance with the present invention;

Figure 8 illustrates, in block diagram form, a framer receiving a transport stream signal;

Figure 9 illustrates, in timing diagram form, the relationship among individual data signals comprising a transport stream;

Figure 10 illustrates, in flow form, a state diagram for implementing a function associated with the framer of Figure 3;

Figure 11 illustrates an algorithmic state machine associated with the framer of Figure 3;

Figure 12 illustrates, in tabular form, global status registers associated with a portion of Figure 6;

Figure 13 illustrates, in tabular form, interrupt registers associated with a portion of Figure 6;

Figure 14 illustrates, in tabular form, global control registers associated with a portion of Figure 6;

Figure 15 illustrates, in block and logic form, a portion of a framer in accordance with the present invention;

Figure 16 illustrates, in block and logic form, a transport packet parser in greater detail;

Figure 17 illustrates, in block and tabular form, a data output controller in greater detail;

Figure 18 illustrates, in tabular form, video control registers associated with a portion of Figure 6;

Figure 19 illustrates, in tabular form, auxiliary PID control registers associated with a portion of Figure 6;

Figure 20 illustrates, in flow diagram form, a method in accordance with the present inventions;

Figure 21 illustrates, in block and logic form, a video packetized elementary stream parser in greater detail;

Figure 22 illustrates, in flow diagram form, a method in accordance with the present inventions;

Figure 23 illustrates, in block diagram form a video packetized elementary stream parser;

Figure 24 illustrates, in tabular form, global status registers associated with a portion of Figure 6 and fully associated with Figures 21 and 23;

Figure 25 illustrates, in tabular form, interrupt registers associated with a portion of Figure 6 and fully associated with Figures 21 and 23;

Figure 26 illustrates, in block form, an output controller and memory in accordance with the present invention;

Figure 27 illustrates, in flow diagram form, a method in accordance with the present invention;

Figure 28 illustrates, in block diagram form, a detailed view of an adaptation field parser;

Figure 29 illustrates, in tabular form, global status registers associated with a portion of Figure 6 and fully associated with Figure 28;

Figure 30 illustrates, in tabular form, interrupt registers associated with a portion of Figure 6 and fully associated with Figure 28.

Figure 31 illustrates, in tabular form, global status registers associated with a portion of Figure 6 and fully associated with Figure 28;

Figure 32 illustrates, in block diagram form, an alternate embodiment of a transport packet demultiplexor;

Figure 33 illustrates, in block diagram form, a detailed view of a private packet packetizer of figure 32;

Figure 34 illustrates, in block form, representations of private packets from the packetizer of figure 33;

Figures 35-38 illustrate, in flow diagram form, a method of splicing video in accordance with the present invention;

Figures 39-42 illustrate, in flow diagram form, a method of performing blind acquisition of an MPEG-2 data stream; and

Figure 43 illustrates, in block form, a general purpose computer system in accordance with the present inventions.

Figure 48 illustrates, in block diagram form, an alternate embodiment of a portion of the system of Figure 7;

Figure 49 illustrates, in block diagram form, an alternate embodiment of a portion of the transport packet parser of Figure 16;

Figure 50 illustrates, in block diagram form, an alternative embodiment of a portion of the PESP of Figure 23;

Figure 51 illustrates, in flow diagram form, a method in accordance with a specific embodiment of the present invention;

Figures 52 and 53 illustrate, in tabular form, portions of the register set accordance with the system of Figure 7.

Detailed Description of the Drawings

Sw A1
In accordance with a specific embodiment of the present invention, detection and/or handling of error condition is enabled for an error being in the transport packet stream. A determination is made whether received data contains detectable error. If so, an error recovery operation is performed that includes the following: notification of the host CPU, MPEG video decoder, a display processor, and/or disregarding the data. The present invention is better understood by example with reference to the Figures herein, and in particular with reference to Figures 48-53

In operation, the TS core 400 receives transport stream packets. Each packet is synchronized to the TS core 400, and demultiplexed. Each packet is demultiplexed based upon its Packet Identifier (PID), which identifies the type of data carried in the packet. The TS core 400 is bufferless in that no packet data is stored within the TS core 400 for access by video or system processing. Instead, the demultiplexed data is stored in one or more locations within each of the Video memory 471, and the system memory 472.

Transport Stream Core 400 includes a Framer 410, Transport Packet Parser 420 (TPP), a PES Parser (PESP) 430, Adaptation Field Parser (AFP) 450, Buffer Controller 460, and register set 480.

The register set 480 is further illustrated in Figure 6. Generally, the register set 480 includes interrupt mask registers, control registers, status registers, and system interface registers. Interrupt mask registers are used to enable or disable specific interrupts. Control registers specify how various aspects of the TS core 400 are to operate. Further examples of types of control registers include Global Control Registers; Video Control Registers, which control how video packets are handled by the TS core; and Non-Video Control Registers, which control how non-video packets are handled by the TS core.

In operation, the framer 410 receives a raw transport stream which is analyzed to isolate and provide individual transport stream packets (TSP) to the bus 405. In one embodiment, the bus 405 receives byte wide data (the data bus width could also be 16 or 32 bits) and a control signal to indicate when the current byte of data is valid. In addition, the Framer 410 generates a signal labeled PACKET START to indicate the first byte of a packet, and a signal labeled IN SYNC to indicate when the data on the bus 405 is synchronized, or locked onto by the Framer 410.

108E31
The TPP 420 is connected to the bus 450, and receives the IN SYNC and PACKET START signals. Parsing of a TSP (packet) by the TPP 420 is enabled when the IN SYNC signal and the PACKET START signals are asserted indicating the beginning of a new packet. During parsing of the header portion of a packet the PID number is obtained. Based upon the value of the PID number, registers are updated, and a determination is made whether the TSP is to be saved, further processed, or discarded.

When it is determined to save the packet, the TPP 420 asserts the signal labeled EN_TPP which is received by the Buffer Controller 460. Based upon this enable signal, the Buffer controller 460 retrieves the packet data and stores it in a predefined memory location.

When it is determined to discard the packet, no further action by the TPP 420 is needed, resulting in the remainder of the TSP being ignored.

1284

When it is determined to further process the packet by one of the other parsers 450 or 430, the TPP 420 asserts one of their respective enable signals. For example, if it is determined that the packet contains video data, the TPP 420 will assert the signal labeled PESP EN, likewise, if it is determined that the packet contains adaptation field data, the TPP 420 will assert the signal labeled AFP EN. Based upon these signal being active, the respective parser will further process the packed data.

In response to being enabled by the TPP, the Video PES Parser 430 further processes the packet by parsing the header of the video PES. Based upon information carried in the header of the video PES, registers are updated, and the video payload may be stored or discarded.

1285

When it is determined to save the video payload, the PESP 420 asserts the signal labeled EN_PESP which is received by the Buffer Controller 460. Based upon this enable signal, the Buffer controller 460 retrieves the packet data and stores it in a predefined location of video memory.

The Buffer controller 460 receives and stores the data payload based upon control signals received from the parsers. Because the packet data is stored directly in the system memory 472, associated with a main system (not shown), or the video memory 471, associated with a video adapter (not shown), the packet data is not stored in TS core 400. Therefore, the core 400 and each of its parsers are described as bufferless. By storing data directly in the system memory 472 and the video memory 471, the system does not have to access memory space within the TS core 400. This eliminates delays associated with the prior art which occurred when the system had to wait on TS core bus accesses to be completed before the needed data could be retrieved.

The bus connections between the buffer controller 460 and the system memory 472 can vary depending upon the implementation chosen. For example, both the video memory 471 and system memory 472 can be connected to the buffer controller 460 through a PCI (Peripheral Components Interconnect) bus, or the system memory 472 can be connected to the buffer control 460 through a PCI bus, while the video memory 471 is connected to the buffer controller 460 through an AGP (Accelerated Graphics Port).

Figure 7 illustrates another embodiment of a TS core in accordance with the present invention. The TS core of Figure 7 includes framer 710, TPP 720, AFP 750, PESP 730, buffer controller 750, and registers 780.

The registers 780 are analogous to registers described with reference to Figure 7.

The framer 710 provides transport stream data labeled FRAMER DATA on an eight-bit bus, (or 16 or 32) and provides a signal labeled FRAMER DEN. The FRAMER DATA is an eight-bit wide data byte, or word, which has been received from the transport stream. The FRAMER DATA is qualified by the signal FRAMER DEN, which is an enable signal. The signal FRAMER DEN is asserted during each valid FRAMER DATA.

The FRAMER DATA and FRAMER DEN signals are provided to each of the parsers of Figure 7, and the Buffer controller 760. The TPP parser 720 receives the header information of new packets when the framer 710 asserts an IN SYNC signal and an PACKET START signal. The combination of these signals, when asserted, indicate that the present FRAMER DATA is part of the packet header. As a result, the TPP 720 receives the FRAMER DATA from the data bus for parsing.

In a specific embodiment, the IN SYNC signal provided by the framer 710 will be active whenever the framer 710 is locked onto, or synchronized with, the transport stream. If the IN SYNC signal is deasserted, the TPP will not receive the data. Furthermore, the PACKET START signal can be a single pulse asserted during the first byte of a new packet, or it can be a signal that is asserted during the first byte of the packet and not deasserted until the last byte of the packet. The first byte of the packet can be defined in different manners. For example, the first byte can be defined to be the sync byte of a packet, or the first byte after the sync byte.

Based upon the PACKET START signal, the TPP 720 can maintain a byte count indicating the location of a current byte within the packet being received. Based upon this count, the TPP 720 will parse the header of the packet which is the first four bytes.

During parsing of the packet header, the TPP receives the PID of the current packet. Based upon the PID value, the TPP can enable other parsers to perform additional parsing operations. For example, when the PESP 730 of Figure 7 is a dedicated video PES parser, and the PID associated with a packet received by the TPP is the video PID, the TPP will enable the PESP 730 by asserting the signal labeled VIDEO. Additionally, TPP asserts the signal labeled VSTART when the current frame is the first frame of a PES stream. This indicates to the PESP that the elementary stream header is at least partially within the current frame. The VSTART signal allows the PESP to locate and parse the header of the video PES, while the VIDEO signal allows subsequent video payload to be retrieved. Likewise, when the adaptation field control of a packet header indicates that adaptation field data is to follow, the TPP will provide a signal labeled AFSTART to indicate the beginning of the adaptation field. In response, the AFP 750 will parse the adaptation field of the current packet.

When a current packet, that is not a video packet, is to be received by the TS Core of Figure 7, the TPP will receive the packet from FRAMER DATA and provide the entire packet one byte at a time as TPP DATA to the Buffer controller 760. Similarly, when the packet is a video packet, the PESP 730 will receive video data payload from the FRAMER DATA and provide it to the Buffer controller 760, which is subsequently framing data bytes into double words as PESP DATA. Any data associated with the adaptation field of the packet will be provided to the buffer controller 760 from the AFP parser 750 as AFP data.

In response to the various data and control signals received from the parsers, the buffer controller stores the data. In a specific mode, where all packets are to be stored, the FRAMER DATA and control signal FRAMER DEN can be received directly at the buffer controller 750 for storage.

In accordance with the present invention, each of the parser modules 720, 730, and 750, and the framer 710, as well as any other module which may be included, are implemented to have modular layouts. For example, the layout of the TPP 720 is modular when its layout is performed independent of the layout of any of the other

The TSTART signal, when provided, is used to indicate when transmission of a new transport stream packet occurs. When TSTART is available, the synchronization

process is trivial because the provider of the transport stream NIM is required to specify the start of each new packet.

The ERROR signal, when present, indicates that the data being received may be corrupted due to a potential error in the data path. ERROR the decoder that the information at this point is at best suspect, if not incorrect.

As previously indicated, various combinations of signals comprising the transport stream can occur. In addition, the format of individual signals can vary. For example, TCLOCK can qualify the TDATA signal on either a rising edge or a falling edge. In accordance with a specific embodiment of the present invention, the TCLOCK edge that qualifies TDATA can be defined in the framer 710.

Another transport stream variation is how the TDATA is transmitted to the framer 710. TDATA can be transmitted in one of either MSB first or a LSB first mode. When transmitted in MSB first mode the most significant bit of each data byte is received first, and in LSB first mode the least significant bit of each data byte is received first. In accordance with a specific embodiment of the present invention, whether data is transmitted LSB first or MSB first can be defined in the framer 710 to properly receive bytes of TDATA.

Another transport stream variation is the polarity of an active control signal. For example, the control signal can be active at either a logic level 1 or a logic level 0, depending upon the system implementation. In accordance with a specific embodiment of the present invention, the polarity of control signals can be defined in the framer 710 to properly identify the correct asserted logic level.

TDATA can be received bit-by-bit, byte-by-byte, or by other various word size. Within the received stream, the individual units of data are referred to having a location. For example, the first byte associated with the data stream is referred to being at a first location, while the 188th byte would be referred to as the 188th location. The use of the term "location" also implies a point in time, whereby a first byte would occur at a first

time, or period, and the 188th byte would occur at a 188th time period as references to the TCLOCK.

Figure 9 illustrates the relationship between the various control and data signals of the transport stream. Specifically, Figure 9 illustrates a TCLOCK signal having a rising edge for qualifying each data byte of the TDATA signal. Likewise, in the illustration of Figure 9, the TVALID signal is always asserted during the first byte indicating that the data is valid. The TSTART signal is synchronized to the first byte of the TDATA signal, which is a synchronization byte. In a specific implementation, the synchronization byte of the TDATA signal will always have the Hexadecimal value 47h. The ERROR signal is not illustrated, however it would be asserted to indicate when an error has occurred.

While the timing diagram of Figure 9 does not explicitly show bits of TDATA being received serially, it should be understood that for each byte representation of TDATA in Figure 9, 8 individual data bits can be received, qualified by eight TCLOCK pulses, to form the bytes illustrated. When TDATA is received in a bit-by-bit manner, without a TSTART signal, there is no knowledge as to which of the bits being received represents the first bit of a byte, where by "first bit" it is meant the first bit received when the device is turned on and started latching the data. Likewise, the same is true for the first byte, let alone which byte represents the first byte of the frame. The state diagram of Figure 10 is a state diagram describing synchronizing the decoder core 700 to the transport stream being received.

Figure 10 illustrates a state diagram for the framer 710. The state diagram of Figure 10 includes four states. State A is the synchronization lost state. State B is a synchronization search state. State C is a synchronization verify state. State D is a synchronization lock state. Upon a hardware or software reset the system in accordance with the present invention enters State A through transition 504. When in State A, synchronization to the data packets has been lost. When synchronization to the transport stream has been lost, it is not known where a new packet begins or an old packet ends. As a result, it is not possible to receive data. Note that when a TSTART signal is provided as part of the transport stream synchronization is known and guaranteed,

therefore State C, synchronization verify, is will not be entered if TSTART is active. For illustration purposes, this diagram assumes that the incoming stream is already byte aligned and that there is no need to look for byte boundaries.

The first byte is a sync byte for the transport stream, and has a predetermined value. In MPEG-2 implementation, the synchronization byte has the hexadecimal value 47h. Transition path 511 loops into State A whenever a byte received is not equal to the synchronization value 47h. During the transition 511, a synchronization lock counter (SyncLockCnt) is set to a stored value. This value of the synchronization lock counter indicates the number of consecutive successful synchronization bytes that must be detected prior to determining the system is synchronized. In the specific implementation, each byte is received by the framer is compared to the value 47h. In one embodiment where a serial bit-stream is received, and the byte boundary within the bit stream is not known, the bit-stream is monitored on a shifted 8-bit basis in order to monitor every possible combination of the bits in the stream to detect the synchronization value. The transition path 513 is taken once the synchronization value is detected.

During transition 513, the synchronization lock counter is decremented to indicate a successful detection of the synchronization value. By identifying a first synchronization byte, the synchronization lock count is decremented the first time. Note that if the synchronization byte value is equal to 47h and the synchronization lock count is equal to zero the transition 512 is taken to State D to indicate successful synchronization.

From State B, transition path 522 is taken for each received byte until the expected location of the next synchronization byte is reached. Because MPEG-2 transport stream packet is 188 bytes long, there will be additional 187 bytes before the next synchronization byte is expected. This is necessary because the value 47h might occur elsewhere in the stream (i.e. this value is not a reserved value). Therefore, on the 187 byte of the packet transition path 521 is taken to return to State A so that the next byte can be evaluated. If at State A it is determined that the 188th byte is has a valid synchronization value of 47h the state machine will transition either on transition path 512, or transition

path 513 depending on how many valid synchronization bytes have been identified. In the event that the 188th byte doesn't have synchronization value, transition 511 is taken, the synchronization lock count is set to the synchronization lock register value, and the system returns to State A.

By transitioning in the manner described between State A and State B, the framer 710 is able to monitor a data stream and determine a valid synchronization location using only the TCLOCK and TDATA signals. Once the valid synchronization location has been identified, by receiving a predefined number of correct sync values, the transition path 512 is taken to State D.

State D indicates that the framer 710 has currently obtained a synchronization lock state. However, in order to assure that the data stream continues to be a valid data stream, the transition 542 is used to determine when the next expected sync byte location is to occur. Transition 541 places the system in state C at reception of the transport stream sync byte to verify synchronization. If synchronization is verified, the system transits to state D along transition path 533. As a result of transitioning along path 533, a drop count is reset to a stored value, which indicates how many sync bytes must be missed before synchronization is lost. This way the incoming stream is continuously monitored for any errors.

By allowing for a predefined number of missed synchronization bytes, intermittent glitches can be ignored. This is useful depending upon the desired data integrity. For example, a drop count value of three would indicate that more than three lost synchronization values in a row will result in the state machine entering State A the synchronization lost state.

When the synchronization value is not found, transition path 533 is taken back to state D. As a result, the synchronization drop count (SyncDropCnt) is decremented to indicate the sync value was not valid, but SyncDropCnt is not yet zero. When the synchronization value is not found, transition path 531 is taken to State A when the synchronization drop count is equal to zero indicating synchronization has been lost.

In the manner indicated, the state machine of Figure 10 allows synchronization to be detected by framer 710 based upon a predetermined number of recognized synchronization values. The predetermined number specifies how many valid packets need to be detected sequentially before it is determined that a valid synchronization lock state has occurred. Likewise, once a valid synchronization lock state has been encountered, the number of missed transport stream packets that must occur can be user defined.

Figure 11 illustrates an Algorithmic State Machine (ASM) diagram in accordance with the framer. Upon reset the flow proceeds to step 602.

At step 602 a variable labeled ByteCnt is set equal to zero indicating the current byte is believed to be the transport stream sync byte, while the variable InSync is also set equal to zero indicating the system is not yet synchronized. At step 602, the framer 710 is in a state labeled frame_byte indicating that the current byte is expected to be a transport stream sync byte, and therefore is to be evaluated.

At step 603, a determination is made whether or not a current byte being evaluated is equal to the hexadecimal value of 47h. When not equal to this value, the flow proceeds to step 621. At step 621, a variable synchronization lock count (SyncLockCnt) is set equal to a register value that specifies the number of valid synchronization bytes needed before synchronization is declared. From step 621 flow proceeds back to step 602.

If at step 603 the synchronization byte value is detected, the flow continues to step 604. At step 604 a variable byte boundary (Byte_Boundary) is set equal to a value bit count (BitCnt), which is zero at step 604.

At step 605 the synchronization lock count variable is decremented to indicate a successful transport stream sync frame detection.

At step 606 a next byte is received. At step 606, the framer 710 is in a state labeled sync_search to indicate the next expected sync byte is being identified.

At step 607, a determination is made whether or not the next byte is the byte to be evaluated for synchronization. If the byte is not to be evaluated the flow proceeds to steps 622 and 606 where the byte count is incremented and a new byte received. In this manner the loop comprising step 606, 607 and 622 is expected until the next byte is the expected sync byte to be evaluated is received, and the flow proceeds to step 608.

At step 608 the variable ByteCnt is set equal to zero, allowing for the next transport packet to be identified. Also, the InSync flag is set equal to zero. At step 608, the framer 710 is in a state labeled sync_lost.

At step 609 a determination is made whether or not the current byte has a value equal to the synchronization value. When the value is not equal to the synchronization value a further determination is made at step 623 whether or not the TSTART signal is active. If the TSTART signal is active, indicating that the start of the transport stream is occurring, the flow will proceed to step 608 for further synchronization. However, if the TSTART signal is not active, or not currently used, the flow will proceed to step 602 for further synchronization. If at step 609 a determination is made that the synchronization value matches the current byte, the flow will proceed to step 610.

At step 610, the variable SyncLockCnt is decremented to indicate successful detection of the transport stream sync value.

At step 611 a determination is made whether or not the synchronization lock count value has been met indicating the framer has locked onto the transport stream. In the specific example, since the synchronization lock count is decremented, when the SyncLockCnt value is equal to zero the condition has been met. If the desired number of consecutive synchronization bytes have not been received, the flow proceeds to step 608. However, if the desired number of consecutive synchronization has been made the flow proceeds to step 612.

At step 612, the synchronization drop count is set equal to the register value indicating how many sync frames must be missed before synchronization is declared lost, and an interrupt is issued to indicate synchronization lock (SyncLock) has been obtained.

At step 635, a variable InSync is set equal to one to indicate the system is synchronized to the transport stream. Therefore, at step 602, the framer is in a state labeled sync_lock.

At step 636, a determination is made as to whether or the current byte is the expected sync byte value. If not, the flow proceeds to steps 644 and 635 receiving the next byte and incrementing the byte count value. If so, the flow proceeds to step 632. At step 632 the InSync variable is maintained equal to one, and the byte count variable is set to zero. At step 632, the framer is in a state labeled sync_verify.

At step 633 a comparison is made of the value of the received byte in order to determine if it is equal to the synchronization value. In the event the byte does match the synchronization value flow proceeds to step 634, where the sync drop count register is set equal to a predefined register value. By setting the sync drop counter value equal to the register value, it is indicated that a predefined number of synchronization values must be missed before synchronization is deemed to be lost.

If at step 633 the synchronization value is not encountered, the flow continues at step 641. At step 641, the SyncDropCnt is decremented to monitor how many synchronization bytes missed.

At step 642, a determination is made whether synchronization has been lost. Specifically, synchronization has been lost if SyncDropCnt is equal to zero. If so the flow will continue at step 643. If not, the flow continues at step 635 previously discussed.

At step 643 the SyncLockCnt is set to the number of a valid synchronization values which must be recognized before synchronization lock is achieved, and an interrupt is generated indicating that synchronization has been lost (SyncLost). The flow proceeds from step 643 to step 624.

At step 624, a determination is made whether or not the signal TSTART is active. In the event TSTART is not active the flow will proceed to step 602 in the manner

previously discussed. In the event that the TSTART is active the flow will proceed to step 608 where the proper synchronization signal will be monitored.

One skilled in the art will recognize that the state diagram of Figure 10 and the ASM diagram of Figure 11 implement similar methodologies in order to accomplish synchronization to a transport stream using the framer 710.

Figures 12-14 illustrates specific registers capable of being utilized to implement specific framer features. For example, various variables described herein are described in the registers of Figures 12-14

Figure 12 illustrates the status and state registers of the framer 710. A field labeled FramerSyncLock is used to indicate that frame synchronization has been acquired, this is analogous to State D of Figure 10, and/or having arrived at sync_lock, step 635, of Figure 11.

A field labeled FramerSyncDrop is utilized to indicate when synchronization has been lost. This is analogous to State A of Figure 10, and/or having arrived at SyncLost, step 608, of Figure 11. This is analogous to the FramerSyncLock variable.

The register Field labeled CurrentFramerState indicates one of five states. In a first state, the framer is in the process of capturing a new byte. In a second state the framer is out of transport packet frame synchronization. In the third state, the framer is searching for synchronization. In a fourth state of the framer is checking for synchronization. Finally, in the third state, the framer is in transport packet frame synchronization. Depending upon the location within the state machine of Figure 10, or the diagram of Figure 12, the values of Figure 12 will be updated.

Figure 13 illustrates a list of the interrupt registers. A field labeled enable global demultiplexer interrupt (EnableGlobalDemuxInterrupt) is utilized to enable all core interrupts. When negated all the core interrupts would be disabled.

An event interrupt mask field (EventInterruptMask) is utilized to mask specific interrupt sources including the FrameSyncLock interrupt, and the FrameSyncDrop

interrupt. The framer synchronization drop bit is used to disable an interrupt that would occur when drop synchronization drop has occurred.

Figure 14 illustrates a portion of a configuration register illustrating various field options associated with the framer. A framer sync lock length field (FramerSyncLockLength) comprises 5 bit field used to select the number of consecutive transport packets, with valid sync bytes, that need to occur sequentially to determine synchronization lock has occurred. The field FramerSyncLockLenth is analogous to the variable SyncLockReg of Figures 5, and the register value indicated at steps 621 and 643 of Figure 11.

A framer sync drop length field (FramerSyncDropLength) comprises a 3 bit field to select a number of consecutive transport packets that must be consecutively missed before the synchronization is declared lost. The field FramerSyncDropLenth is analogous to the variable SyncDropReg of Figures 10, and the register value indicated at steps 612 and 634 of Figure 11.

A framer bit polarity field (FramerBitPolarity) is a single bit used to indicate whether the transport stream data is being received MSB first or LSB first.

A framer clock polarity field (FramerClockPolarity) is a single bit that when asserted indicates transport stream data that is latched on a rising clock edge. Conversely, when negated, data is latched on a falling clock edge.

A framer mode field (FramerMode) comprises two bits for defining a combination of external transport stream control signals to be used to determine synchronization. In a first mode, only the TSTART value is used to determine when the latched data is valid. In a second mode, the TVALID signal is used determine when synchronization is valid. In the third mode, the framer will use both TSTART and TVALID to determine when synchronization is valid. In the fourth mode, the framer will use TCLOCK and TDATA to latch the bit stream in.

Each of the control signals TVALID, TSTART, and TERROR have an respective register fields TVALID Polarity, TSTART Polarity, and TERROR Polarity to indicate whether the signals are active high signals, or active low signals.

By providing the control information described in the configuration registers of Figure 14, it is possible to program a decoder core 700 in order to interface to a large number of transport stream protocols.

Figure 15 illustrates a specific implementation of a portion of the framer 710 using the control register information. The implementation utilizes various configuration registers to select modes of operation. In the specific embodiment illustrated, the transport stream data (t_data) is received serially and loaded into four registers 1010 through 1013. The serially loaded data is provided at a parallel output associated with each of the registers. The parallel outputs of registers 1010 and 1011 are received at inputs of multiplexer 1020. Parallel outputs of registers 1012 and 1013 are provided to the inputs of a multiplexer 1021. The parallel outputs from the multiplexers 1020 and 1021 are provided to inputs associated with the multiplexer 1022. The output of multiplexer 1022 is provided to two bit shifters 1030 and 1031. Parallel outputs of the bit shifters 1030 and 1031 are provided to a comparator 1040.

In operation, registers 1010 and 1011 receive the serial bits of data on rising clock edge, while registers 1012 and 1013 receive the serial bits of data on falling clock edge. This assures proper storage of data without knowledge of TDATA's relationship to TCLOCK. Clock triggers registers 1010 and 1011 store the data either from left-to-right, or from right-to-left. By loading data from opposite directions it is assured that whether data is received MSB first or LSB first that the data is stored in a manner consistent with the architecture. For example, a hexadecimal value 11h stored in register 1010 will be stored as a hexadecimal value of 88h in register 1011.

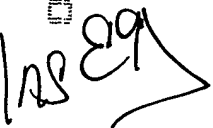
Register field FramerBitPolarity is to select either the MSB first registers 1011 and 1013, or LSB first registers 1010 and 1012, while the register field FramerClockPolarity will select the register having the appropriate clock qualification.

The data provided to the bit shift registers 1030 and 1031 is shifted bit-by-bit, to provide all possible byte combinations to the sync byte comparator 1040, which determines when the synchronization value has been encountered, and asserts the control bit in response to a successful compare. When a successful compare occurs, it is assumed that the byte and Packet boundaries have been located.

The synchronization hardware illustrated below multiplexer 1022 of Figure 15 is isolated from the external clock. This is advantageous over the prior art, in that a loss of the TCLOCK signal does not shut down the control logic associated with synchronization of the transport stream.

In accordance with the present invention, it is possible to provide a flexible framer capable receiving a variety of physical transport stream formats and determining synchronization when only clock and data are present, and to provide appropriate synchronization control signal.

One skilled in the art will recognize that many specific implementations of the framer can be incorporated. For example, the framer may have a first in first out (FIFO), or other type buffer associated with it. In addition, instead of selecting specific configuration parameters using registers, other configuration specification means could be used, such as making them pin selectable, or any other of various types methods capable of describing selectable features.

128 89  ~~Figure 16 illustrates a more detailed view of the TPP 620. TPP 620 further includes storage locations 721, a counter controller 722, register controller 723, video PID location 724, and adaptation field start detect circuit 725.~~

In the implementation shown, each of the various header fields of a transport stream packet have a storage location within the storage locations 721. Because the transport stream data is received on a byte-by-byte basis, and because the header field locations are fixed, the data for the individual fields is readily obtained. In the embodiment of Figure 16, each storage location for a specific data field is connected to

THE **NEW** **YORK** **PUBLIC** **LIBRARY**

[illegible][illegible][illegible][illegible]

in the allocation table 727. Therefore, each valid entry of the allocation table 727 is compared to the current PID value stored in storage location 721. If any of the valid entries match, the Output data controller 726 will provide the entire packet to controller 760 for storage. Because the current PID value is not available until after the fourth byte of the header is received, the output data controller 726 saves the first three byte in case they need to be stored.

The allocation table 727 illustrated lists 32 PID indexes (PID_0-PID_31) which have PID values associated with them. The allocation table 727 can actually be an abstraction of register locations from the register set 780. Figure 18 illustrates video control registers, which are a portion of the register set 780. The PID value associated with the PID_0 entry of the allocation table 727 is defined to be the active video PID value, which is received from the VideoPID field of the Video Control Registers of Figure 18. Likewise, Figure 19 illustrates Demultiplexer Control Registers, which are a portion of the register set 780 used to identify packets, other than current video packets, which are to be saved. The PID values associated with the PID_1 through PID_31 entries of the allocation table 727 are received from their respective register locations within the Demultiplexer Control Registers of Figure 19. For an entry to be valid, the EnableParsing field of the PID register needs to be enabled.

If a received packet's PID number is not listed in the PID Allocation Table, the packet is not processed further, i.e. discarded, and the next received TSP is analyzed. However, if the PID of the current packet is listed in the PID allocation table, and it is not the video PID, the packet is saved to memory.

Figure 20 illustrates a method associated with the TPP parser. At step 211, the TSP is received by the framer as discussed with reference to Figures 10 and 11 herein.

At step 212, the packet is received at the TPP. In the manner discussed herein, the packet is made available to the TPP one byte at a time. The framer provides an indicator where the first byte of the packet is located.

In response to receiving the first byte of the packet, the TPP will parse the packet header at step 213. From the parsing of the header, the TPP will retrieve the PID value of the packet.

At step 214, a determination is made whether the packet is identified as a valid packet. As previously described, one way to be identified as a valid packet is be specified in an allocation table, which can be derived from specific register information. When a PID value is listed in the allocation table, the packet is to be further processed.

At step 215, a determination is made whether the packet is a packet that is to be additionally parsed. For example, step 215 specifically indicates that a determination is being made whether the PID value indicates the packet is a video packet. If so, flow proceeds to step 226 for video parsing as indicated in Figure 22. If the PID does not indicate a packet for special processing, i.e. not a video packet, the flow proceeds to step 227 where the data is send the buffer controller for storage, as indicated with reference to Figure 22.

When the PID allocation table, or other means, indicates the packet is a video packet the Packetized Elementary Stream Parser (PESP) is enabled to allow further processing. In the specific embodiment of the PID allocation table listed above, the video PID is stored as PID_0. However, other methods of identifying the video PID, such as the use of a flag or other indicator are also possible. The operation of the PESP is controlled by the PESP Control Registers, as illustrated in Figure 18.

Figure 21 illustrates the PESP 730 in greater detail. PESP, of Figure 21, includes a counter controller 752, storage location 751, register controller 753, data output controller 756, video data control module 755, and portions of register set 780.

In the implementation of Figure 21, a storage location within the storage locations 751 is reserved for the STREAM ID header field of a transport stream packet. In the embodiment shown at Figure 21, inputs of the storage location for STREAM ID are connected to the appropriate bits of the data bus and the counter controller 752, to receive stream ID data from the FRAMER DATA representation of the transport stream at the

correct time. The counter controller 752 receives the VSTART signal indicating the start of a new video PES and generates enable signals to capture the stream ID, and other information, from the video PES header. The counter controller is controlled by the signals VIDEO, FRAMER DEN, and VSTART. The VIDEO signal indicates the current packet is a video packet. The FRAMER DEN signal indicates when the current FRAME DATA byte is valid, and VSTART indicates when the current packet is the first packet for the video PES, in other words, VSTART indicates when the video packet contains video PES header data to be parsed. Based upon the VSTART signal, and the FRAMER DEN signal, the counter controller 752 can determine which byte of the header is currently being received.

1ns 210
In another implementation, control module 755 is controlled by the EnableParsing field (not shown in Figure 21) of the video control registers of Figure 18. The EnableParsing field is a one bit field, which when deasserted, prevents further parsing of the video packet by the video PESP. Therefore, when the EnableParsing field is negated, the header of the video packet would not be parsed, and therefore, the packet would be discarded. The counter controller can be controlled directly from the EnableParsing bit of the video control registers, or indirectly where the VIDEO signal is disabled by the TPP 620 in response to the EnableParsing bit being deasserted.

Once the video PES header field has been parsed, register fields dependent upon the specific fields of the video PES header can be updated. The register set 780 is accessed and updated by the register controller 753 of Figure 21, which is connected to storage locations 751 of the PESP. In addition, the register controller 753 can retrieve or access register data as needed. For example, the values EnableParser, ProcessStreamID, and StreamID are register values from register set 780.

The video data control module 755 contains logic that enables the video data payload of the present packet to be stored. Operation of the control module 755 is determined in the content of various video control registers, as illustrated in Figure 18. The EnableParsing field is a one bit field, which when negated prevents any data from the current video packet from being saved. The ProcessStreamID field is 1 bit-field. When

asserted, it enables further parsing based upon a specific stream ID value of the video PES header, such that if the video control register field StreamID of Figure 18, does not match the parsed steam ID from the current packet, the data of the packet will not be saved. This is an advantage over the prior art, where filtering on the stream ID field of the video PES was done in software, generally by the system.

In the specific implementation illustrated, only the data payload of the video PES is stored. Since the parsing is done in hardware, there is no need for the header information to be stored.

In another embodiment, the field labeled StartFromPUSICommand is used to indicate whether video PES parsing is to start immediately with the next packet or wait until a new PES stream is received as indicated by the VSTART signal, where the acronym PUSI stands for Payload Unit Start Indicator and is a part of MPEG-2 syntax. Once the new video PES stream is identified, the StartFromPUSICommand field is negated, and all subsequent video packets are further processed by the PESP.

The video PESP parser is bufferless in that no local buffers are used to store the payload data for access by other parts of the system. The prior art parsers stored the parsed data in large buffers locally, which were then capable of being accessed by system components by requesting access tot he local bus. The bufferless parsers of the present invention do not store data locally for access by the system. Instead, parsed data to be buffered is transmitted to the buffer controller 460, which buffers data in system or video memory.

Figure 22 illustrates a method associated with a video PESP parser. At step 216, the PESP has received an indication that a video packet is ready to be parsed. The notification can be directly from the TPP, by a polling mechanism, or other type interrupt. Step 16 determines whether parsing of the video stream is enabled. This can be determined based upon the field labeled EnableParsing of the video control registers of Figure 18. When parsing of the video packet is not enabled, a specific action will occur. One action would be to perform no further processing of the packet, as illustrated. In another implementation, the packet would be automatically stored without further

24
am7

parsing, perhaps with the packet header field. When parsing of the video packet is enabled, the flow proceeds to step 217.

At step 217, a determination is made whether the packet is to be parsed immediately, or whether parsing of video packets is to wait until a new video PES is detected. If the packet is to be parsed immediately, the flow proceeds to step 223. If the packet is not to be parsed immediately, flow proceeds to step 218 to determine when the proper criteria for parsing is met. Field StartFromPUSICommand indicates whether parsing is to be immediate.

At step 218, a determination is made whether the present packet is the first packet of the video PES. If the packet is a new video PES packet, the field StartFromPUSICommand is disabled, and flow proceeds to step 219. If the new packet is not the first packet of a video PES, the flow will terminate as indicated with no further processing.

At step 219, a determination is made whether the current video packet is to be parsed based upon the packet stream ID. If so, flow proceeds to step 220, if not, flow proceeds directly to step 223.

At step 220, the PESP parses the stream ID from the PES header as discussed with reference to Figure 21. In addition, Figure 23 illustrates addition hardware parsing which can be performed by the PESP.

At step 222, a determination is made whether the parsed steam ID from the PES header is equal to the value stored at register field StreamID of the video control registers of Figure 18. If so, the field StartFromPUSICommand is disabled to allow subsequent packets associated with the video PES to be stored, and flow proceeds to step 223. If no match occurs, the flow terminates as indicated, and no further processing occurs.

Ins E12

At step 223, the packet data is sent to the buffer controller for storage, as discussed with reference to Figure 25.

Note that additional parsing steps can occur between steps 217 and 223, such that from step 217 additional parsing would occur. The parsing steps illustrated in Figure 22 are all by-passed if the current transport stream packet is to be immediately routed to a system memory and parsed by the host processor.

Figure 23 illustrates additional parsing features of the PESP 730. Figure 23, includes a counter controller 752, storage location 751, register controller 753, and data output controller 756.

In the implementation of Figure 23, a storage locations within the storage locations 751 are reserved for the specific PES header field of the Packetized Elementary Stream . In the embodiment of Figure 21, inputs to the storage locations 751 for specific PES header fields are connected to the appropriate bits of the data bus, while the counter controller 752, which receives the VSTART signal indicating the start of a new video PES, receives PES header data from the FRAMER DATA representation of the transport stream at the correct time.

The counter controller 752 will generate enable signals to capture the various PES header fields based upon the values stored in locations 736, and a counter value generated by counter 737. The counter controller is controlled by the signals VIDEO, FRAMER DEN, and VSTART. The VIDEO signal indicates the current packet is a video packet. The FRAMER DEN signal indicates when the current FRAME DATA byte is valid, and VSTART indicates when the current packet is the first packet for the video PES, in other words, VSTART indicates when the video packet contains video PES header data to be parsed. Based upon the VSTART signal, and the FRAMER DEN signal, the counter controller 752 can determine which byte of the header is currently being received. As indicated with reference to the StartFromPUSICommand register of Figure 18, the counter controller can either allow for immediate PES parsing upon receiving a video packet, or it can wait to parse the PES data until a packet containing PES header information is received. Where PES parsing is immediate, the video data is provided to the output buffer.

In operation, a compare operation determines if the present counter 737 values is equal to any of the values stored in location 736. If so, it indicates that the current clock cycle is providing data to be stored in one of the fields of storage locations 751. As a result, the controller 752 will generate an enable to the appropriate one or more fields represented in the current clock cycle, and the field data will be latched.

The compare function 738 can be implemented in many different ways. For example, a state machine or logic can be used to indicated which of the storage locations 751 are to be stored at a specific time. In addition, feedback is provided to the controller 752 from various storage locations 751 to assure proper operation. For example, all PES header will have the field portions 766 of storage location 751. However, depending upon various values of these, and other fields, the field portions 767-769 may or may not be present in a particular PES header.

For example, whether the fields portions 767 exist in a current header is determined by whether the binary framing indicator "10" immediately follows the PES packet length field as in Figure 3, and as stored in the OptionalHeader indicator field. This OptionalHeader indicator field is compared to the expected value and the results are provided to the controller 752 to indicate additional parsing is to be done. As a result, the parser 752 continues to generate control signals to store the fields associated in the field portions 767.

In a similar manner, the Flags field of storage portion 767 indicates which of the storage portions 768 are present, and the ExtentionFlags of storage portion 768 indicate which of the storage portions 769 are present. In this manner, the controller 752 determines which header fields are present and need to be stored in storage locations 751.

Once the video PES header field has been parsed, register fields dependent upon the specific fields of the video PES header can be updated. The register set 780 is accessed and updated by the register controller 753, which is connected to storage locations 751 of the PESP. In addition, the register controller 753 can retrieve or access register data as needed. Figure 24 illustrates a subset of the Status Register Fields

associated with on implementation the PESP, while Figure 25 illustrates interrupt mask registers having corresponding bits.

Once the header has been completely parsed, the data associated with the payload portion of the current PES packet can be provided to the data output controller 756 as discussed with reference to Figure 21. In an alternate embodiment, the 16 bytes of optional PESPrivate data associated with the PES header and stored in storage locations 769 are provided external the PESP to a private data packetizer as will be discussed in greater detail herein.

Figure 26 includes a detailed view of buffer controller 460 of Figure 5, video bus/memory controller 488, system bus/memory controller 468, video memory 471, and system memory 472. In the specific embodiment illustrated in Figure 26, the buffer controller 460 includes a data path for handling video PES data to be stored in the video buffer 500 of video memory 471, and a data path for handling other PES data that is to be stored in system memory buffers 501-503 of the system memory 472. The data path for handling other PES data includes the System FIFO (First In First Out) controller 466, FIFO 462, and System HBI (Host Bus Interface) Controller 463. The data path for handling video PES data includes a Video FIFO controller 486, FIFO 461, Video HBI Controller 483. Each of the System and Video data paths accesses transport demultiplexer register 465.

13213
In operation, the System FIFO controller 466 provides an interface between the Parsers of Figures 5 and 7 and the FIFO 462. The controller 460 allows filtered packet data to be received and stored in the FIFO 462. Once stored in the FIFO 462, the System HBI controller 463 requests access to the video memory 471 through the controller 468. The controller 468 may include a system bus controller, a memory controller, or a combination of a memory/system bus controller. Generally, the controller 468 will control access to other system resources as well.

In accordance with the invention, the System Memory 472 has been partitioned by the system host to include one or more system circular buffers 501-503. The system buffers 501-503 are implemented as circular buffers and are filled by operation of the

controller 483. The controller 483 handles the buffer “write” pointer. The “read” pointer for the buffers is managed by the software on the system host side (not shown) which retrieves data from the buffers 501- 503. There can also be a “high water” mark register associated with each buffer (not shown). The purpose of a “high water” mark register is to provide an interrupt when the write pointer crosses the value in this register. However, because there is generally only one interrupt for each of the plurality of buffers, software polling can be used to determine the cause of the interrupt.

In a specific implementation, the number of system buffers is limited to 15 buffers. The transport core may use fewer than 15 buffers. More than one PID per buffer is allowed. However they have to be different, i.e. the same PIDs can not be allocated to more than one buffer (i.e. one TSP packets can be routed into only one destination ring buffer). The Transport Demultiplexer registers of Figure 19 are used to specify where data associated with a specific PID is to be stored. For each PID to be saved, a buffer index is used to specify one of the 15 buffer locations in system memory. Multiple PID types can be stored at a common buffer by specifying the same buffer in the BufferIndex field. In an alternate embodiment, data associated with all system PIDs can be stored to a single buffer, which may be specified or defined by default. Note with reference to Figure 7, the buffer index, or location, can be determined by one of the parsers, and provided to the Buffer controller 760.

The video data path of Figure 26 is handled in a manner analogous to the system path described above. However, in the specific embodiment, only one buffer, in video memory, is associated with the video path.

The physical memory location and the size of the ring memory buffers 500-503 is specified by the system host using buffer configuration and management registers (not shown). The host processor has to initially specify buffer start address and length of the buffer. Other buffer data can also be used, for example, a threshold register can be implemented.

The size of the video buffer depends on horizontal and vertical pixel resolution, frame rate, profile and level, maximum bit rate and video buffering verifier (VBV).

ATSC requires a buffer of minimally 0.95MB (VBV = 488); while for MPEG-2 Main Level at High Profile, the size is 1.17MB (VBV= 597). The buffer controller 460 will manage a write pointer for the video stream. The read pointer is managed by the control 488 associated with the video adapter. Hardware or software can generate an interrupt if the write pointer is equal to the read pointer – 1 (overflow condition).

Regarding the audio buffer requirements, the worst case is for a when the nominal audio bit rate 640 kbps with sampling frequency of 32 kHz. The actual size of the compressed bit stream audio buffer depends on a priority and the rate of occurrence of the audio decoder thread, when audio is decoded in software.

On a channel change, software will flush the buffers by setting the read pointers to be equal to the write pointers after the transport stream parser has been turned off.

108214
09/23/00 10:50:00
Figure 27 illustrates a method in accordance with the present invention describing the operation of the system HBI controller 463 of Figure 26. The flow is also applicable to the video HBI controller 463. At step 801, a determination is made whether there is data stored in the FIFO 462. If not, flow remains at step 801 until data is present, otherwise, the flow proceeds to step 810. At step 810, the buffer to which the data is to be stored is identified. The destination buffer is identified when matching and crossing the PID number, or identifier, to the buffer number in the transport demultiplexer register 465. The buffer can be identified by accessing the allocation table, or by receiving a buffer index from the transport parser or other portion of the transport core.

At step 802, a determination is made whether the identified buffer is full, or otherwise not capable of receives additional data. If the buffer is not capable of receiving additional data, the flow loops back to step 802 through step 811, which implements a delay. Note the delay of step 811 may be a fixed delay, as result of polling to determine if the buffer is full, or the delay of step 811 may be variable, such as where the delay is based an interrupt which indicates when the buffer is available. Once the desired buffer is no longer full, flow proceeds to step 803.

At step 803, a bus request is made to allow access to one of the buffers 501-503. Once the bus connected to the buffer has been acquired, the next block of data is written to the appropriate buffer at step 804. A block of data can be a word, double word, or any other size of data specified by the system. In a specific embodiment, the parsers of Figure 5 assure data is provided to the FIFO only in whole blocks by always writing entire blocks of data to the FIFOs.

At step 805, a determination is made whether the identified buffer 501-503 is now full. If so, the flow proceeds to step 807 where the bus is released, if not full, the flow proceeds to step 806, where it is determined if more data is to be written.

At step 806, a determination is made whether more data resides in the FIFO 462. If so, flow proceeds to step 804, otherwise, the flow proceeds to step 807 where the bus accessing the Buffer is released and flow returns to step 801. In another embodiment, the bus would be released after each block is transferred, instead of determining if more data is to be written. By implementing the flow of Figure 26, the data stored in the FIFO 462 is transferred to the buffers.

128215
The buffer implementation described provides an advantage over the prior art in that moving the buffers in to system and video memory associated with an external system, such as a general purpose computer, allows for bufferless parsers. As a result, the system and video resources do not have to wait to access buffers local to the parsers. The performance improvements using bufferless parsers has be observed by the present inventors to be up to 40% over the prior art. In addition, by allowing for parsing of the PES data, it is possible to limit the amount of bandwidth used to store unused data. One skilled in the art will recognize the present invention has been described with reference to a specific embodiment, and that other implementations and variations of the present invention would be anticipated. For example, when a TSP is "sent" from the TP to the PESP or the buffer controller, it is to be understood that not necessarily all of the header information need be sent. In fact, in would generally be necessary for only the PID associated with the packet be forwarded. In addition, the location and implementation of

the register sets and functionality described herein can be partitioned in ways other than the specific implementations described.

The AFP parser 750, illustrated in Figures 5 and 7, parses data associated with the adaptation field of a transport packet. The Transport Packet Parser 720 enables operation of the Adaptation Field Parser 750 when the adaptation field of the header indicates the presence of an adaptation field. Figure 28 illustrates, in block diagram form, a detailed view of the Adaptation Field Parser 750.

The AFP 750 illustrated in Figure 28 includes adaptation control counter 782, latch 785, register logic 786, and storage and register locations 781, 783, and 784. In operation, the adaptation control counter 782 receives signals on connections labeled AF START, FRAMER DEN, and FRAME DATA. The connection labeled AF START receives signals from the Transport Packet Parser 720, and indicates the beginning of the transport packet's adaptation field. The connection labeled FRAMER DEN receives signals from the Framer 71, and indicates when each new byte of data is available on the FRAMER DATA bus. Based upon the received signals, the adaptation control counter 782 provides the control signals necessary to parse specific field information from data received on the FRAME DATA bus.

In operation, the Transport Packet Parser 720 will assert a signal on to the connection AF START in response to the adaptation field control portion of the transport packet header indicating the presence of an adaptation field. The signal on the AF START connection will be asserted in relation to the assertion of the first byte of adaptation field data onto the FRAMER DATA bus.

The first byte of the adaptation field indicates the length adaptation field. Therefore, the adaptation control counter 782 will latch the first byte of the frame data into a storage location labeled AF LENGTH to determine the length of the adaptation field. Accordingly, the adaptation field has a variable length between 1 and 183 bytes long. By decrementing the adaptation field length by one as each new byte of data is received on the FRAME DATA bus, the adaptation control counter 782 can monitor which fields, or field portions, of the adaptation field are present on the FRAME DATA

bus at a specific time. Accordingly, the adaptation control counter 782 provides operational control signals to each of the storage locations of storage portion 781 to correspond to the presence gets data on the FRAME DATA bus.

Generally, the storage locations 781 correspond to specific registers of the register set 780 of Figure 7. For example, the discontinuity indicator field is known to be the first bit, of the second byte, of the adaptation field. Therefore, the storage location labeled Discontinuity Indicator of storage area 781 will be connected to only the first bit of the FRAME DATA bus. Furthermore, logic associated with the Adaptation Control Counter 782 will provide a latch control signal to the Discontinuity Indicator of storage location 781 only when the counter associated with the Adaptation Control Counter 782 indicates the second byte of data is present on the FRAME DATA bus. In a similar fashion, the other specific adaptation bit-fields associated with locations 781 will be parsed. Note that the individual locations of storage locations 781, may be the actual register locations of register set 780, or may be storage locations local to the adaptation field parser 750. Where the storage locations are local to the parser 750, a register control portion (not shown) can be used to update values within the register set 780.

The Optional Flags field of storage locations 781 is connected back to the adaptation control counter 782 in order to provide feedback. The need for this feedback is best understood with reference to prior art Figure 1. Prior art Figure 1 illustrates that the adaptation field can include five optional fields. The presence of each of these five optional fields is indicated by flag bit. The field labeled Optional Flags represents the five flags of the adaptation field to indicate the presence of the optional fields. Therefore, by providing feedback from the optional field location, the adaptation control counter 782 can correctly determine which optional field data is present, and when to receive optional field data.

Storage locations 784 generally represent register locations whose values are determined based upon the parsed information of storage locations 781. Specifically, the register logic 786 monitors the operation of the adaptation control counter 782 and/or the contents of the locations 781 to determine when the PCR value has been received. In

When the optional flags indicate that the optional fields includes transport private data, the adaptation control counter 782 will provide the private data from the FRAME DATA bus to a PESP PRIVATE DATA bus through either directly, or through a latch 785. In addition, the adaptation control counter 782 will provide a signal to a node labeled PRIVATE DATA ENABLE to indicate when the PESP PRIVATE DATA bus has valid data. In one embodiment, the PRIVATE DATA ENABLE node is clocked for each valid byte of data written to the PESP PRIVATE DATA bus. In another embodiment, the PRIVATE DATA ENABLE node would include multiple nodes, whereby one node pulsed to indicate each valid byte of data written to the PESP PRIVATE DATA bus, while the other node indicated the valid PESP private data cycle. The valid PESP private data node would be asserted for the entire assertion of PESP private data from a common transport packet.

Figure 29 illustrates a number of status register fields associated with the register set 780 of Figure 7, which are associated with the operation of the adaptation field parser 750. Video AFPcrReceived is a single read bit, which is set to 1 after arrival and extraction of the PCR sample in the adaptation field. The assertion of this bit will cause an interrupt be generated if the VideoAFPPcrReceived bit of the event interrupt mask is enabled. Subsequent read access of this field will cause it to be cleared.

40

Register field VideoAFDiscontinuityFlag is a single bit R/W field that is set to 1 after a discontinuity indicator flag has been asserted in the adaptation field of the video transport packet. Assertion of the discontinuity indicator flag indicates discontinuity on the continuity counter. The assertion of this bit will cause an interrupt to be generated if the VideoAFDiscontinuityFlag of the event interrupt mask register of Figure 30 is asserted. The subsequent access of this field by software will cause this field to be cleared.

Register field VideoAFRandomAccess is a single bit R/W field that is set to 1 when the video packet has a random access flag asserted in the adaptation field. This indicates the start of the new elementary stream. The assertion of this bit will cause an interrupt be generated if the VideoAFRandomAccess bit of the event interrupt mask register of Figure 30 is asserted. The subsequent access of this field by software will cause the field be cleared.

Register field VideoAFSplicingFlag is a single bit R/W field that is set to 1 when the video packet has the splicing point flag asserted in the adaptation field. This flag indicates that a splicing point is approaching. The assertion of this bit will cause an interrupt to be generated if the VideoAFSplicingFlag bit of the event interrupt mask register of Figure 30 is asserted. The subsequent access of this field by software will cause the field to be cleared.

Register field VideoAFSplicingPoint is a single bit R/W field that is set to 1 when the video packet has the VideoAFSplicingFlag asserted and the AF Splice Countdown register has a value of 0. The setting of this bit is controlled by the register logic 786 of Figure 28. The assertion of this bit will cause an interrupt to be generated if the VideoAFSplicingPoint bit of the event interrupt mask register of Figure 30 is asserted. The subsequent access of this field by software will cause the field to be cleared.

Register field VideoAFPrivateData is a single bit R/W field, which when set to 1 indicates the video packet has adaptation field private data. The assertion of this bit will cause an interrupt be generated if the Video AF Private Data bit of the event interrupt mask register of Figure 30 is asserted. The subsequent access of this field by software

will cause the field to be cleared.

Register field AFSpliceCountdown is a byte wide R/W field that contains the current splice countdown value from the current adaptation field.

ins 87
~~Figure 31 illustrates control registers associated with register set 781 that control operations associated with the Adaptation Field Parser 750.~~

Register field EnabledAFPrivateData is a single bit R/W field that when asserted enables parser of the adaptation field private data.

Register field AFPrivateDataBufferIndex is a four bits field which specifies 1 of up to 15 destination buffers in the system memory where the private data is to be stored.

Register field PCRIndex is a five bits field which specifies one of 32 PID values from which the PCR data is to be extracted.

Register field Enabled Auto Splicing is a single bit R/W field that enables automatic splicing of the video elementary stream.

Figure 32 illustrates an alternate embodiment of a portion of the transport demultiplexer illustrated in Figure 7. Specifically, Figure 32 illustrates a Private Data Packetizer 740 connected to the Adaptation Field Parser 750 and the Video PESP 730. The Adaptation Field Parser 750 is connected to the Private Data Packetizer 740 through the AFP PRIVATE DATA bus and node PRIVATE DATA ENABLE. The Video PESP 730 is connected to the Private Data Packetizer 740 through the bus labeled VIDEO PRIVATE DATA and node labeled VIDEO PRIVATE DATA ENABLE. The Private Data Packetizer 740 receives private data from the AFP 750, and video PESP 730 and associated control signals. In turn, private data packetizer 740 provides the private data packet on a bus labeled PRIVATE DATA to buffer controller 760, and a control signal on the node labeled PRIVATE DATA ENABLE to the buffer controller 760.

In operation, private data associated with the video PESP 730 has a fixed length of 16 bytes. However, the private data associated with the transport packet, which is parsed by the AFP 750, has a variable length from 1 to 181 bytes. Because the FIFOs

761 and 762 of the buffer controller 760 store double words, it is possible, and in fact likely, that the private data associated with the adaptation field of transport packet will not provide to the FIFOs private data that ends on a double words boundary. The significance of this is best understood with a discussion of the operation of one embodiment of the buffer controller 760.

During normal operation of the buffer controller 760, bytes of data associated with transport packets from the Transport Packet Parser 720 and video data from the Video PESP 730 are provided to the buffer controller 760. Each parser has a corresponding double word buffer in the buffer controller 760, which receives and stores the individual bytes of data until an entire double word has been received. For example, the first byte of data provided by the Transport Packet Parser 720, is stored in the first byte location of a double word buffer 763, while the second, third, and fourth data bytes will be stored in second, third, and fourth byte locations of the double word buffer 763 respectively. When the double word buffer 763 has received the fourth byte, the double word is written from the double word buffer 763 to the FIFOs 761, thereby freeing up the double word buffer 763 for subsequent bytes.

When a specific data packet of a packetized elementary stream does not end on a double word boundary, the double word buffer 763 will be partially filled and therefore not send the end of the reception of the specific data packet. However, this is not a problem if specific data packet is repeatable over many packetized elementary streams, because the subsequent data packet associated with the same packetized elementary stream can be expected to be received within a relatively short amount of time to complete filling the double word buffer. Once the double word buffer is filled using data bytes from the subsequent data packet, the field double word buffer 763 will be sent to the FIFOs 761.

While it is not problematic for a specific packet of a packetized elementary stream to not completely fill the double word buffers associated with the buffer controller 760, the same is not true of private data associated with specific transport stream or packetized elementary stream. This is because the private data associated with packetized

elementary stream has a fixed length and is not streaming data of the type associated with the transport packet parser 720 or the video of the video PESP 730. Because the private data from be transport packet has a variable length, there is no guarantee that the private data will end on a double word boundary. If the private data does not end on a double word boundary, the partial double word portion of private data at the end will not be sent to the FIFO until additional private data from unrelated source is received. Therefore, the system software that interprets private data, would have incomplete data. The private data packetizer 740, illustrated in Figure 32, addresses this problem.

In operation, the private data packetizer 740, can receive private data from each of the Adaptation Field Parser 750, and the Video PESP Parser 730, and forms a private data packet to be sent to the buffer controller 760, which is guaranteed to have a length that ends on a double word boundary. Note that both the packetized elementary stream data from the Transport Packet Parser 720, and the private data generated by the private data packetizer 740 are sent to the system FIFOs 761. An index indicator, which specifies which circular buffer in system memory the private data is to be stored, is provided to the FIFOs 761. The index indicator is specified in the register field labeled AFPrivateDataBufferIndex, which was discussed with reference Figure 31 herein. Therefore, all private data, whether from the Adaptation Field Parser 750 or from the Video PESP 730, is been provided to the same buffer in system memory. Figure 33 illustrates the private data packetizer 740 in greater detail.

The private data packetizer 740 of Figure 33 includes counter controller 741, the AF Data Type Code storage location 742, PES Data Type Code storage location 743, Stuffing Code storage location 744, AFP Data Latch 745 PESP Data Latch 746, and fixed length Indicator Code 747.

INSE 181

~~The AF Data Type Code storage location 741 stores the specific eight-bit type indicator associated with the adaptation field private data. The PES Data Type Code storage location 743 stores the specific the eight-bit type indicator associated with the PESP private data. The Stuffing Code storage location 744 stores the specific eight-bit stuffing code which is used to pad private data packet to insure the private data packet~~

218
m4

always ends on a double word boundary. The AFP Data Latch 745 is used to store the actual private data from the adaptation field parser to be provided to the buffer controller 760. Similarly, the PESP Data Latch is used to store the actual private data from the PESP parser is to be provided to the buffer controller 760. The fixed length indicator code 747 stores the fixed length value associated with the PESP parser private data. In the specific example, the PESP parser private data will always be 16 bytes of data, or 0x10 hexadecimal.

004310 6969450

In operation, the counter controller 741 can be enabled either by be AFP Data Enable signal, or by the PESP Data Enable signal. When the counter controller 741 is enabled by the PESP Data Enable signal, the number of bytes of private data is fixed at 16 bytes, therefore, the value of 16 will be used by be AFP counter controller 741 to generate the appropriate signal for the PRIVATE DATA bus. Unlike PESP private data, AFP private data has a variable length. The actual number of bytes of AFP private data, not including the length byte, is transmitted in the first byte of the AFP private data field of the data packet. Therefore, the counter controller receives the number of bytes of transport packet private data by latching the first byte of the private data field of the data packet. The first byte of the private data field is received on the AFP DATA bus on or after the PES Data Enable signal has been received. Based upon the source of the private data and the length of private data, the private data packetizer 740 will construct the private data packet.

Figure 34 illustrates a specific embodiment of a private data packet generated by the private data packetizer 740. Data block 771 illustrates the format of the private data packet having specific fields: type, length, data, and stuffing.

The type field of the private data packet indicates the source of the private data, either transport packet private data, or video PES data. In a specific embodiment, the hexadecimal value of 0x55 is used to indicate private data associated with the transport packet received from the AFP 750, and a hexadecimal value of 0xAA is used to indicate private data from the video PES received from the video PESP. Note that in other embodiments of the present invention, additional data types can be received from other

sources.

The length field of the private data packet specifies the length of private data that is to follow the private data packet. Note that in the specific embodiment illustrated, the value of the length field does not include a count for the length field byte.

The stuffing field of the private data packet is used to assure the private data packet ends on a double word boundary. As indicated, the stuffing field can include zero to three bytes of data.

Data block 772 of Figure 34 illustrates a private data packet associated with the video PES stream. Specifically, private data packet 772 has a type value of hexadecimal value 0xAA indicating the private data packet is associated with a video PES. The length field of packet 772 has a hexadecimal value of 0x10, which indicates that 16 bytes are contained within the subsequent data field. Accordingly, the subsequent data field of the private data packet includes 16 bytes, 0 through F. Because the length of the fields type, length, and data is not one, the number of stuffing bytes needed to insure the private data packet ends on a double word boundary is readily determined. Therefore, two stuffing bytes of 0xFF are represented in the stuffing field of private data packet 772.

By adding to stuffing bytes in the stuffing field of the private data packet, the length of the private data packet ends on a double word boundary. Therefore, when the data bytes of the private data packet 772 of Figure 34 are provided to the buffer controller 760, and more specifically to the double word buffer 764 of the buffer controller 760, it is assured that the entire private data packet will be provided to the FIFOs 761 without delay. Once the data is provided to the FIFO 761, the double word of data will be provided to the appropriate buffer in system memory. System software will, therefore, be able to access the private data stored in system memory without delay, and determine the type of data based upon type field, and length of the data based upon length field. Private data packet 773 of Figure 34 illustrates another specific private data packet. Specifically, packet 773 has a type value of hexadecimal value 0x55, which represents an adaptation field data packet. The length field of the packet 773 has a hexadecimal value of 0x0F, which indicates 15 bytes of data are associated with the adaptation field private data. As

a result, 15 bytes of data 0 through E are represented in the data field, and a total of three stuffing bytes are needed in order to assure that private data packet ends on a double word boundary.

Data blocks 774 and 775 indicate other specific private data packets associated with the transport packet provided by the adaptation field parser 750. Specifically, the length of the private data has been varied in packet 774 and 775 to illustrate packets having a single stuffing byte, and no stuffing bytes respectively.

Referring to Figure 33, the counter controller 741 provides the appropriate data type code to the buffer controller 760 by selecting storage location 742 or 743 respectively depending upon whether adaptation field private data or video PES private data is being received. As a result, the appropriate data type will be provided onto the bus labeled PRIVATE DATA. Note the bus labeled PRIVATE DATA is illustrated to be a wired-OR bus, however any type of multiplexing would be the appropriate.

The length field of the private data packet is provided to the PRIVATE DATA box differently depending upon whether adaptation field private data or video PES private data is being provided. The length of video PES private data, which is always 16 bytes, is provided to the PRIVATE DATA bus by selecting the storage location 747, which contains the hexadecimal value 0x10. Enabling the storage location 747 allows the hexadecimal value 0x10 be provided to be PRIVATE DATA bus. The length of adaptation field private data is provided to the PRIVATE DATA bus by latching the first byte of the adaptation field private data into be AFP data latch 745. Because the first byte of the adaptation field private data specifies the number of private data bytes that follow, the appropriate length value for the length field is provided to the PRIVATE DATA bus.

Once the type and length field data have been provided to the PRIVATE DATA bus, the actual data is provided to the PRIVATE DATA bus. This is accomplished in a similar manner for both the adaptation field private data and a video PES private data. Specifically, the counter controller 741 latches the data into one of the appropriate data latches AFP data latch 745, or PES data latch 746. In response, be associated private

data is provided to the PRIVATE DATA bus. Note the private data could be provided to the PRIVATE DATA bus directly to transmission gates, or any other appropriate logical interface.

The generation of stuffing codes, from the stuffing code register 744, is controlled by the control counter 741. Because the control counter 741 knows the length of the private data provided, it can readily determined the number of bytes needed, if any, to assure the private data packet ends on a double word boundary is readily calculated. Therefore, after the last byte of the private data, from either the AFP or PESP, the appropriate number of stuffing codes are been provided to the PRIVATE DATA bus by selecting the storage location the 744 determined number of times this assures the buffer controller 760 will receive a number of bytes that and on a double word boundary. As a result, the private data can be provided to the system buffer controller 760 without delay.

Figures 35 through 38 illustrate a method of performing automatic splicing using the data parsed herein provided to their respective registers. The term splicing refers to the process of switching between two video streams. For example, splicing can be used to switch between the video of the main program and a video of a commercial, between video to a commercial, and from a commercial video back to the main program video. The method of Figures 35 through 38 is contingent upon being enabled by the compound to field of the control register. Splice points can be sorted as Out Points or In Points.

An Out Point is a location in the stream where the underlying elementary stream is well constrained for a clean exit (usually after I or P frame). An Out Point Packet is the packet of the PID stream preceding an Out Point. MPEG-2 syntax defines Out Points at transport layer as:

splicing_point_flag =1, splice_countdown = 0, seamless_splice_flag =1;

An In Point is a location where the stream is well constrained for a clean entry (usually just before a sequence_header and I frame at the closed GOP. MPEG-2 syntax defines In Points at transport and PES layers as:

splicing_point_flag=1, splice_countdown = -1, seamless_splice_flag=1;

payload_unit_start_indicator = 1 random_access_indicator = 1, data_alignment_indicator = 1;

At step 301, of Figures 35, the interrupt mask register is written to, in order to enable reception of interrupts based upon the video splicing flag and the video splicing point. The video splicing flag and the video splicing point values are determined by parsing performed by the Adaptation Field Parser 750. The video splicing flag indicator is one of the optional flags of storage location 781 of Figure 28, and is represented by register field VideoAFPSplicingFlag in the global status register of Figure 29. The video splicing point is determined by register logic 786 of Figure 28, and results in the register field labeled VideoAFSplicingPoint being set when the video splicing flag is set and the splicing countdown register, labeled AFSpliceCountdown, is equal to 0.

At step 302, the method of Figure 35 waits for an interrupt to occur.

At step 303, an interrupt has been received, and a determination is made as to the interrupt type. If the interrupt type is a splice point, the flow proceeds to connector A, which is continued at Figure 38. If the interrupt type is a splice flag, the flow proceeds to step 304. If the interrupt type is a different type of interrupt, the flow returns back to step 302.

At step 304, a determination is made as to the type of splice represented by the splice flag interrupt. This can be determined by analyzing the splice countdown value received by the adaptation field parser 750 of Figure 28. Specifically, if the splice countdown value is a positive value it is an indication that the splice flag has identified an out-splice point. An out-splice point indicates that the current video elementary stream

being received is about to end, and the flow proceeds to connector B, which continues at Figure 36.

If at step 304 the splice countdown value is equal to zero, it is an indication the splice flag has identified a splice point. The splice point is as identified as that point in time where video is to be switched from a current video stream to a next video stream. A splice point flag is set and the adaptation field parser 750 of Figure 28, when the splice flag is asserted, and the splice countdown register is equal to 0. (Note that under normal operation, the splice point path from step 304 will not be taken because the splice point should have been detected at step 303, thereby bypassing step 304).

If at step 304, it is determined that the splice countdown value is negative, it is an indication that an in-splice point has been identified. And in-splice point indicates that the current video elementary stream being played is just began, and the flow proceeds to connector C that continues in Figure 37.

One skilled in the art will recognize that specific register values identifying splice-in points and splice-out points could be provided in the same manner the separate register location was provided for the splicing point. Likewise, many other variations of the specific flow herein can be made without departing from the inventive intent.

When out point has been detected at step 304, the flow proceeds to Figure 36. At step 310, of Figure 36, the splice flag interrupt is disabled. The splice flag interrupt is disabled because specific method illustrated in Figures 35 through 38 only needs to execute and out point routine one time. Since the splice countdown value for an out point includes the values from 7 to 1 the out point routine disables the splice flag interrupt at step 310 in order to avoid having to unnecessarily process interrupts caused by the subsequent out point values.

At step 311, splice point interrupts are enabled. Note that splice point interrupts should already be enabled from step 301 of Figure 35. Therefore, the step 311, under ideal operating conditions, will be redundant and not strictly necessary.

At step 312, acquisition of the PAT table is requested. The PAT table is

discussed with reference to prior art Figure 2.

At step 313, a determination is made whether or not the PAT table version number has changed. If the PAT table version number was not changed, the flow proceeds to step 314. However, if the PAT table version number was changed, the flow proceeds to step 317.

At step 317, if the PAT table version number has changed, a new PMT table, see Figure 2, is fetched and the flow proceeds to step 314.

At step 314, a determination is made whether or not the current splice is valid. A valid splice point is recognized by asserted (set to 1) `splicing_point_flag` and `seamless_splice_flag`. Flags are stored in the global status register. If it is determined that the current splice point is not valid, the flow proceeds to step 302 of Figure 35. However, if it is determined that the current splice point is valid, the flow proceeds to step 315.

At step 315, a request is made to receive the break duration time as an optional bit-field available in the `splice_info_section` that indicates an approximate time when the break will be over and when the network In Point will occur. At step 316, the new PID number, received from the new PMT table, is written to a register that shadows the VideoPid register of Figure 18. Referring to Figure 16, the VIDEO PID storage location 424 provides the PID value which identifies the current video stream, while the shadow register associated with location 424 (not illustrated) stores the PID value of the next video stream to be accessed at the splice point. Subsequent to step 316, flow proceeds back to step 302.

When it has been determined at either step 303 or step 304 that the splice point has occurred, the flow proceeds to Figure 38. At step 331 of Figure 38, the PID value stored in the shadow register is transferred into the active PID register. As result, the value stored in the VIDEO PID location 424 of Figure 16 is updated to the new PID value, resulting in the new video PID packets been identified, and selected, as the current video stream. In effect, the newly selected video image will be displayed.

At step 332, a request is made to update the STC counter. The STC counter is updated in order to synchronize the system counter with the new program counter, i.e. a new time base.

1/18/91
At step 333, the splice flag interrupt is enabled. The Splice Flag Interrupt Enable Bit is asserted in order to allow for the recognition of the splice in point. From step 333, the flow proceeds to step 302 of Figure Note that in another embodiment of the present invention, a determination step could be made at the beginning of the flow of Figure 38 as to whether the new PID is associated with a desirable program. If not, an alternate flow ignoring the PID, or using a dummy or alternate PID, could be used. For example, this feature could be used eliminate viewing commercials or other program types.

At step 304, Figure 35, it is determined that an in-splice point is occurring, the flow proceeds to Figure 37. At step 336, of Figure 37, a determination is made whether or not this is the first in-splice point interrupt service request. The first in-splice point interrupt service request, is generally associated with the value minus 1 of the splice countdown register. However, in order to accommodate for possible lost packets, the determination of step 320 may be used along with an indicator as to whether or not to the previous in-splice point interrupt service request has occurred. If this is not the first in-splice point, the flow proceeds to step 302 of Figure 35. If this is the first in-splice point, the flow proceeds step 337.

1/18/91
At step 337, a determination is made whether or not the in-splice point indicator is valid. The in-splice point indicator is validated by determining whether or not random access flag register is set along with discontinuity flag register. The random access flag register, and discontinuity flag register, should both be set because the first packet of a new data stream will indicate the current packet is capable of being randomly accesses by the system, and since no previous packets are associated with the PES stream the discontinuity flag should be set.

At step 338, a request for PMT table acquisition is made. This is done to verify that current PID assignment for a present program is as before the brake (or before added commercial). At step 339, a determination is made whether the PCR and video PID

values are valid. PIDs are verified by analyzing a content of received PMT table for known MPEG-2 program_number. Change formidable if the PCR and PIDs are okay, flow proceeds to step 302 of Figure 35. Otherwise, flow proceeds to step 340.

At step 340, the new PID values are updated.

The method described 32 through 35 provides advantages over the prior art by allowing for the automated handling of in-splice point. Utilizing register values updated by the hardware parsers described herein, automatic splicing is enabled. The amount of saved system software bandwidth provides an advantage over the prior art.

Therefore, one skilled in the art will recognize that providing for hardware parsing of adaptation fields, and the generation of the private data packet regardless of the source of private data, provides advantages over the prior art. In addition, the splicing method described herein, provides for automatic hardware splicing control, which provides advantages over prior art methods of software control splicing.

In another embodiment of the disclosure, a transport stream demultiplexer core register set is initialized to indicate a possible set of transport stream characteristics. An acquisition routine is executed. If acquisition of the transport stream signal does not occur during a predetermined amount of time, the acquisition is not successful. When not successful, the register set is initialized to indicate a different possible set of transport stream characteristics, and the acquisition routine is once executed. This process continues, until the transport stream core acquires the transport stream signal.

Figures 39-42 illustrate a specific implementation of a method for blind synchronization to a transport stream. Blind synchronization allows the framer to acquire the transport stream, i.e. lock onto the transport stream, without any prior knowledge of the transport stream characteristics.

As discussed with reference to Figures 8 and 9, the transport stream can include a variety of signals. At a minimum, the transport stream will include a data signal (TDATA) and a clock signal (TCLOCK). Additional signals that may exist include

TSTART, TVALID, and TERROR. Based upon these signals, the transport stream has a number of characteristics, such as individual signal polarities, and data ordering.

One transport stream characteristic is the polarity of the control signals, which can vary based upon the service provider implementation. In other words, each of the control signals TVALID, TSTART, and TERROR can be either active high or active low signal. As illustrated in Figure 14, each of the control registers has respective register field (T_ValidPolarity, T_StartPolarity, and T_ErrorPolarity) to specify the active edge of each control signal.

An additional characteristic is the data ordering, or the bit polarity. Because the data is received bit at a time, or by bytes, whether the LSB is first or the MSB is first can vary. A field labeled FramerBitPolarity is used to select between a LSB and MSB bit polarity of TDATA.

Another transport stream characteristic is whether the TCLOCK signal latches data on a rising edge or on a falling edge. A field labeled FramerClockPolarity is used to select between these two characteristics.

At step 911 of Figure 39, these variables are initialized to represent a specific set of transport stream characteristics. As further illustrated in step 921 of Figure 40, the registers T_ErrorPolarity, T_StartPolarity, and T_ValidPolarity are set equal to zero. For purposes of discussion, a value of zero will represent an active low polarity, while a value of one will represent an active high polarity.

A variable BIT_ORDER, which corresponds to field FramerBitPolarity, is set to LSB to indicate the LSB of TDATA is to be received first, or right justified bytes of data are received. The variable BIT_ORDER can also be set equal to MSB to indicate the MSB of TDATA is to be received first, or will be left justified where bytes of data are received.

A variable labeled CLOCK_POLARITY, which corresponds to the field labeled FramerClockPolarity in Figure 14, is set to zero, where zero indicates that TDATA is latched on a falling edge.

At step 922 of Figure 40, other initialization overhead functions are performed. For example, Figure 14 illustrates a field labeled FramerMode that specifies the signals associated with the transport stream. Step 912 can include initialization of this field as well.

As step 912 of Figure 39, a verification routine is executed. The verification routine is illustrated in greater detail with respect to Figure 41. At step 931 of Figure 41, reception of the transport stream is enabled. In effect, it implements settings of the current transport stream characteristic. Upon enabling reception of the transport stream, the framer portion of transport stream demultiplexer core begins operation as described previously in an attempt to achieve synchronization.

At step 932 a predetermined amount of delay time occurs to allow the framer to detect a synchronization byte. When the data stream being received is an MPEG-2 data stream the synchronization byte is a hexadecimal 47 (47h). The predetermined delay is used to detect one 188 byte long MPEG-2 packet, and depends on a stream bit-rate and is typically under 100 microseconds.

At step 933, a determination is made whether the synchronization byte was detected. If not, the flow proceeds to step 935, if so, the flow proceeds to step 934.

At step 934, a determination is made whether or not additional synchronization bytes need to be detected before synchronization is declared. In Figure 14, the variable labeled FramerSyncLockLength indicates how many consecutive transport packet synchronization frames need to be detected before synchronization is declared. Based upon this value, the flow will return to step 932 until the specified number of synchronization values has been detected. When the specified number of synchronization frames has been detected, the flow returns to Figure 39 and indicates a successful verification.

At step 935, it has been determined that the transport stream reception has not been successful. As a result, no further attempt is made to acquire the transport stream with the present settings of the transport stream physical characteristic. Therefore, step

935 performs any overhead functions needed, for example, reception of the transport stream can be disabled. Note in other embodiments, reception of the transport stream with improper characteristic settings continues.

From step 935, the verification flow of Figure 41 returns to Figure 39 and indicates that the verification was not successful.

118 E21
~~At step 912 of Figure 39, a determination is made whether the verification routine was successful. If so, the flow proceeds to step 914, if not, the flow proceeds to step 915.~~

At step 914, the proper set of transport stream characteristics has been found and any necessary cleanup occurs. Examples of necessary cleanup items would include notifying the use of successful synchronization, storing of the synchronization characteristics.

At step 915 the transport stream characteristics are incremented. Figure 42 illustrates one method of incrementing the characteristics specified with respect to step 911.

118 E22
~~Figure 42 illustrates a flow diagram that can that increments the transport stream characteristic in such a manner that all possible combinations are covered. By executing this routine, a successful increment will be indicated for all values, except for when BIT_ORDER variable is equal to MSB, and all other characteristics are equal to one. This state indicates that all possibilities have been tested, and an unsuccessful return occurs.~~

At step 916 of Figure 39, a determination is made whether the increment of the transport stream characteristic variables was successful. If so, the flow returns to step 912 where the verification routine is executed again. If the increment of the transport stream characteristic was not successful, indicating that no lock was obtained, the flow proceeds to step 914 for appropriate cleanup.

The present method provides a fast method for acquiring a transport stream having unknown characteristics. Variations of the method described herein would

include varying the number of consecutive synchronization byte required for acquisition to be successful, varying the order in which the variables are varied. Varying the framer mode to indicate the possible combinations of transport stream signals, i.e. clock and data only.

It should be understood that the specific steps indicated in the methods herein may be implemented in hardware and/or software associated with the specific parsers or controller described. For example, a specific step may be performed using software and/or firmware executed on one or more a processing modules. In general, a system for handling transport stream data may include a more generic processing module and memory performing the functionality described herein. Such a processing module can be a single processing device or a plurality of processing devices. Such a processing device may be a microprocessor, microcontroller, digital processor, micro computer, a portion of the central processing unit, a state machine, logic circuitry, and/or any device that manipulates the transport stream. The manipulation of the transport stream is generally based upon operational instructions. The memory may be a single memory device or a plurality of memory devices. Such a memory device may be a read only memory a random access memory a floppy disk memory, magnetic tape memory, erasable memory, a portion of a system memory, and/or any device that stores operational instructions in a digital format. Note that when the processing module implements one or more of its functions to be a state machine or logic circuitry, the memory storing in the corresponding operational instructions is embedded within the circuitry comprising the state machine and/or other logic circuitry.

Figure 43 illustrates, in block diagram form, a processing device in the form of a personal computer system 1000. The computer system 1000 is illustrated to include a central processing unit 1010, which may be a conventional proprietary data processor, memory including random access memory 1012, read only memory 1014, input output adapter 1022, a user interface adapter 1020, a communications interface adapter 1024, and a multimedia controller 1026. Note the central processing unit 1010, the communications interface adapter 1024, and video/graphics controller can individually be considered processing devices.

15E23
The input-output (I/O) adapter 1026 is further connected to, and controls, disk drives 1047, printer 1045, ~~removable storage devices 1046~~, as well as other standard and proprietary I/O devices.

The user interface adapter 1020 can be considered to be a specialized I/O adapter. The adapter 1020 is illustrated as connected to a mouse 1040, and a keyboard 1041. In addition, the user interface adapter 1020 may be connected to other devices capable of providing various types of user control, such as touch screen devices.

The communications interface adapter 1024 is connected to a bridge 1050 such as is associated with a local or a wide area network, and a modem 1051. By connecting the system bus 1002 to various communication devices, external access to information can be obtained.

The multimedia controller 1026 will generally include a video graphics controller capable of displaying images upon the monitor 1060, as well as providing audio to external components (not illustrated).

In accordance with the present invention, the transport core can be implemented at various portions of the system 1000. For example, the transport core can be part of the Communication Interface Adapter 1024, as part of the Multi-Media Controller 1026, or as a separate component of the system connected directly to the bus 1002. In a specific embodiment, the video memory of Figure 5 resides within the Multi-Media Controller 1026, while the system buffers 501 to 503 would generally reside in RAM 1012. In other implementations, a unified memory can be used.

Verification of a system implementing various hardware and methods described herein can be accomplished by implementing an MPEG-2 transport demultiplexer in C language in order to verify the parsing transport files to the level of elementary streams. The produced elementary streams have been compared against streams generated by reference transport stream generation equipment, which in this case, has been used as a reference transport/PES stream decoder.

1. The first step is to identify the key components of the system. This involves understanding the hardware and software involved, as well as the data flow and processing requirements.

- 59

verified with the transport packet header syntax analyzer software. This completes the verification of a C model.

A transport stream generator, or composer is described below for generating test streams, and for testing hardware and software models of transport stream demultiplexers. These streams are generated through a C program and text file containing script command language. It consists of one through several hundred MPEG-2 transport stream packets that would contain different PID's, PES data, etc. Each test stream can scope a specific function of the parser. These test streams are files that can be fed to a behavioral model that reads this file and feed it as a serial bit-stream to the VHDL. These tests can be run to result in bit-maps residing in either the frame buffer memory, or in the system memory .

A data record, in accordance with a specific embodiment of the present invention, is described below capable of implementing the tests used for verifying a C model, or a hardware device or a system. The data record illustrated below allows simple ASCII, or other text based, description of TP or PES syntax.

; This is a comment line

FILE "YourFilename.ES" ; Name of the file containing MPEG-2 elementary stream to be packetized into PES and TP formats.

; The following portion is the PES, also referred to as a data stream, description, and is relevant to PES layer syntax.

[REPEAT PES	M1]	
[PES_SCP	N0]	
[PES_SID	N1]	
[PES_SC	N2]	
[PES_DAI]		
[PES_C]		
[PES_OOC]		
[PES_PTS_FLAG	N3	N4]
[PES_ESCR_FLAG		N5]
[PES_ESRATE_FLAG		N6]
[PES_DSM_FLAG		N7]
[PES_ACI_FLAG		N8]

```
[PES_CRC_FLAG                N9]
[PES_HEADER_DATA_LENGTH      N10]
[PES_PD_FLAG                  [N11]  [RANDOM]  ["TextString1"]]
[PES_PH_FLAG                  N12]
[PES_PPSC_FLAG]
[PES_PSTD_FLAG                N13]
[PES_EXT2_FLAG                N14]
[PES_STUFFING                 N15]
[PES_PAYLOAD_LENGTH  N16 [N17] [RANDOM] ["TextString2"]]
[END]
```

; The following portion is the transport packet description

```
[REPEAT_TP      M2]
[[TP_SYNC N18 N19 [[RANDOM] [BITSHIFT]]]
[TP_TEI N20] [TP_TEI_CONSECUTIVE N21 N22] [TP_RANDOM_TEI]
TP_PID N23
[TP_TSC N24 N25] [TP_TSC_CONSECUTIVE N26 N27 N28]
[TP_RANDOM_TSC N29]
[TP_CC N30 N31 [DISCONTINUITY]] [TP_CC_RANDOM]
[TP_PAYLOAD [N32] [RANDOM] [BITSHIFT] ["TextString3"]]
[[[TP_DUPLICATE N33] [TP_DUPLICATE_RANDOM]] [DOUBLE]]
[END]
```

; The following portion is an adaptation field description

```
[AF_LENGTH      N34]
[AF_DI          N35]
[AF_RAI         N36]
[AF_PCR_FLAG    N37 N38]
[AF_OPCR_FLAG   N39]
[AF_SP_FLAG     N40 N41]
[AF_TPD_FLAG    N42 [N43] [RANDOM] ["TextString4"]]
[AF_EXT_FLAG    N44]
```

All optional fields are given in [] braces; a syntax is relevant to a single PID transport stream. In the specific implementation, multiple PID transport stream files are created by muxing single PID files with simple utility program. All numbers (N1..N43) are given as decimal numbers. Additional syntax relevant to the **PES data record** listed above is as follows:

- M1 (DWORD) specifies a number of PES packets under REPEAT - END structure.

- **PES_SCP** is the packet_start_code_prefix of PES packet set to numeric value N0 (DWORD). If [PES_SCP N0] is not given, packet start code should be set to default 0x000001. Purpose of [PES_SCP N0] is to define error on PES packet start code. N0 is 24bit. If more than one PES packet exist, this error is injected on the first PES packet only, all other PES packet start codes are 0x000001.
- **PES_SID** is PES stream_id field set to numeric value N1 (BYTE). If command [PES_SID N1] is not given, stream_id is encoded as 0xE0. Purpose of [PES_SID N1] is to define error on stream_id field, which occurs on every PES packet.
- **PES_SC** is the PES_scrambling_control bit-field set to a numeric value N2 (BYTE). If [PES_SC N2] command is not specified, PES_scrambling_control should be encoded as 0, indicating non-scrambled (clear) PES payload. Purpose of [PES_SC N2] is to define error or legal case of scrambled PES payload, which ASIC should skip. $N2 \in \{0,1,2,3\}$ and is applied to first PES packet only, all others are generated as clear packets with PES_scrambling_control == 0.
- **PES_DAI** (BOOL) is data_alignment_indicator. If this line is specified in the script file, the data_alignment_indicator of the PES header has to be set to 1 to first PES packet only. If PES packets are creating from ES video, this flag is set when video sequence header is found.
- **PES_C** (BOOL) is copyright bit. If this command is specified, copyright bit of the PES header has to be set to 1, otherwise it is 0. This applies to every PES packet.
- **PES_OOC** (BOOL) is original_or_copy bit. If this line is specified in the script file, original_or_copy bit of the PES header has to be set to 1, otherwise is 0. This applies to every PES packet.
- **PES PTS_FLAG** is the PTS_DTS_flags bit-field specifying existence of the PTS or PTS & DTS values in the PES header. N3 (BYTE) gives encoding type ($N3 \in \{0,1,2,3\}$), and N4 (ULONGLONG) gives exact value of the PTS (only 33 LS bits are meaningful, others has to be set to 0). This line of the script file allows

verification of the PTS extraction circuit of transport demux ASIC. We don't process DTS value, so if N3 is set to 3, DTS can be encoded as 0 or 0x1FFFFFFFF or the same as PTS. If PES packets are creating from ES video, this flag is set when video sequence header is found, otherwise PTS is applied to the first PES packet only.

- **PES_ESCR_FLAG** (BOOL) is ESCR_flag. N5 (ULONGLONG) encodes ESCR_value. Applies to the first PES packet only.
- **PES_ESRATE_FLAG** (BOOL) is ES_rate_flag. N6 (DWORD) encodes ES_rate value. Applies to the first PES packet only.
- **PES_DSM_FLAG** (BOOL) is DSM_trick_mode_flag. N7 (BYTE) encodes a full DSM byte field of the first PES packet only.
- **PES_ACI_FLAG** (BOOL) is additional_copy_info_flag. N8 (BYTE) encodes additional_copy_info of every PES packet.
- **PES_CRC_FLAG** (BOOL) is PES_CRC_flag. N9 (WORD) encodes 16 bit CRC sum. This allows verification of the ASIC's CRC hardware engine. Specified value is injected to second PES packet only, CRC on all others is calculated. First PES packet has no CRC field.
- **PES_HEADER_DATA_LENGTH** is the PES_header_data_length syntax element. N10 (BYTE) gives actual length, which is adjusted with stuffing.
- **PES_PD_FLAG** (BOOL) indicates that 16 private data bytes exist in the header. If N11 (BYTE) is given, private data bytes are set as numeric string which starts from N11 and increases for 1. If RANDOM is defined, private data bytes are set as random bytes. If text string is given, private data bytes are set to ASCII string with 0xFF padding if the string length is less than 16, for example "PES private data". This applies to every PES packet.
- **PES_PH_FLAG** (BOOL) is pack_header_field_flag. N12 (BYTE) gives the length of the pack header. If N12 is set to 13, than a regular 12 byte MPEG-1 pack_header

is inserted, otherwise it is encoded with stuffing bytes (0xFF). N12 should be some small number ($N12 < 20$). This applies to the first PES packet only.

- **PES_PPSC_FLAG** (BOOL) is the `program_packet_sequence_counter_flag`, indicating that every PES packet has PES sequence counter.
- **PES_PSTD_FLAG** (BOOL) is the `P-STD_buffer_flag`. N13 (WORD) encodes `P-STD_buffer_size` (13 LS bits). For parsed video stream, `P-STD_buffer_scale` is always 1, indicating units of 1KB. This applies to every PES packet.
- **PES_EXT2_FLAG** (BOOL) is the `PES_extension_flag_2` indicating second extension field whose length is specified with N14 (BYTE). ASIC skips over 2nd extension field and stuffing area. The second extension is added to the first PES packet only. N14 should be some small number ($N14 < 20$).
- **PES_STUFFING** N15 (BYTE) indicates the number of stuffing bytes added to every PES packet header. $N15 < 32$.
- **PES_PAYLOAD** N16 (WORD) indicates length of the PES payload, whose content can be specified as numerical string starting from value N17 (BYTE) with roll over to 0 after 255, random numerical string or ASCII text string. Only one of the 3 options is valid at the time, for example text string "This is a PES payload" where `PES_PAYLOAD_LENGTH` is set to 21. This applies to every PES packet.

Additional syntax relevant to the **transport packet record** listed in the table above is as follows:

- **M2** (DWORD) specifies a total number of transport packets defined in the REPEAT – END structure. It is relevant only if there is no underneath PES stream (no previous PES syntax).
- Line **TP_SYNC** N18 N19 [RANDOM] [BITSHIFT] is used to specify errors on the `sync_byte`. `Sync_byte`, generated as random unsigned 8 bit integer [RANDOM] or bit-shifted variant of initial sync value 0x47 [BITSHIFT] is set from transport

- **TP_TEI** (BOOL) is transport_error_indicator. It can exist with the specified rate (TP_TEI with rate specified as DWORD N20), it can be generated from packet N21 (DWORD) in all N22 (DWORD) consecutive packets; or it can be randomly generated [TP_RANDOM_TEI]. In the first case, repetition rate is N19, i.e. if N20 is set to 0, TEI occurs only once, at the first packet, if N20 is 1, every packet has TEI error flag, if N20 is 2 every second packet has TEI error and so on. If neither of three specifiers is given, **all packets have no TEI flag set**.
- Line **TP_PID N23** specifies the PID of the current transport stream encoded as WORD with 13 LS bits as actual PID value.
- **transport_scrambling_control** can be set as:

constant with command **TP_TEI_CONSECUTIVE** and set to value **BYTE N26** $\in \{0,1,2,3\}$ from transport packet **N27** (DWORD) in all **N28** (DWORD) next transport packets, and after that set to 0, indicating no scrambling.

If this line is not specified, `transport_scrambling_control` is set to 0 (not scrambled) in **all** transport packets.

- **continuity_counter** can be set as to BYTE value $N30 \in \{0,1,2, \dots 15\}$ starting from packet N31 (DWORD). Optional **DISCONTINUITY** specifier, if exists defines a 'legal' discontinuity on the continuity counter field, set in the adaptation field. This line may exist as **TP_CC_RANDOM** where random transport packets are chosen to add random 4-bit value for continuity counter. If this line is not specified, continuity_counter field should start from 0 with first transport packet, and doesn't have any discontinuity **in all transport packets**.
- **Payload of the transport packet** can be set to a numeric string starting from BYTE value N32 with roll over to 0 after 255, or it can be filled with random bytes (**RANDOM**), bit shifted start codes (**BITSHIFT**, i.e. **0x47 bit shifted**) or set to ASCII text string for example "Transport packet payload" as a good, known payload. If FILE directive is specified, a payload is built from underneath PES packets of ES stream and command **TP_PAYLOAD** is illegal in that case.
- Optional line **TP_DUPLICATE N33** or **TP_DUPLICATE_RANDOM** defines duplicate transport packet as the image of the latest transport packet with duplicate rate in DWORD N33 or randomly generated rate of duplication. Optional **DOUBLE** specifier means that a duplicate packet is duplicated creating three consecutive packets with the same content. If N33 is set to 0, **only first transport packet is duplicated, otherwise every N33th transport packets is duplicated**.

Additional syntax relevant to the **adaptation fields portion** listed in the table above is as follows:

- Line **AF_LENGTH N34** specifies designed length of the adaptation field, including stuffing bytes. BYTE N34 encodes length. The actual length depends on specified flags and fields, and possible stuffing and is determined on packet by packet base. N34 allow us to define length in such a way that influence of errors on the adaptation_field_length can be tested. If specified length N34 is in conflict with the actual, computed adaptation_field_length, then adaptation field should contain all specified fields and flags with no stuffing and adaptation_field_length set to N34. If

specified N34 complies with the possible computed length(s), then N34 should be used with actual length adjusted with stuffing.

- **BOOL AF_DI N35** is used for **discontinuity_indicator**. **DWORD N35** is the repetition rate of the **discontinuity_indicator**. If N35 is set to 0, the **discontinuity_indicator** should exist **only once**, in the first transport packet, otherwise it repeats **on every N35 transport packets (if N35 is 1, discontinuity_indicator is repeated in every TP)**.
- **BOOL AF_RAI N36** is used for **random_access_indicator**. **DWORD N36** is the repetition rate of the **random_access_indicator**. If N36 is set to 0, **random_access_indicator** should exist **only once**, in the first transport packet, otherwise it repeats **on every N36 transport packets (if N36 is 1, random_access_indicator is repeated in every TP)**.
- **BOOL AF_PCR_FLAG N37** N38 stands for **PCR_flag** with initial value encoded in the **LONGLONG N38** and the repetition rate defined by **DWORD N37**. If N37 is set to 0, PCR exists only once, **in the first transport packet, otherwise it repeats every N37 transport packets**. N38 should increase by step $100000 / N35$;
- **BOOL AF_OPCR_FLAG N39** stands for **OPCR_flag** with the repetition rate defined by **DWORD N39**. If N39 is 0, OPCR exists only once, **in the first transport packet, otherwise it repeats every N39 transport packets**. Because our ASIC doesn't extract OPCR value, OPCR can be set to the latest PCR value.
- **BOOL AF_SP_FLAG N40** N41 stands for the **splicing_point_flag**. **BYTE N40** indicates initial **splice_countdown** value and **DWORD N41** is its repetition rate. If N41 is set to zero, **splicing_pont_flag** exists **only once**, in the transport packet with **splice_countdown** set to 0 (splicing point), otherwise it repeats on every N41st transport packets. Term: $N41 > N40$;

- **BOOL AF_TPD_FLAG N42 [N43] [RANDOM] ["TextString4"]** stands for the `transport_private_data_flag`; `transport_private_data_length` is defined by BYTE N42. Data content can be set to increasing numeric value defined by BYTE N43 (step is +1), or to the random byte numeric string, or to the ASCII text string with 0xFF padding to the length N42. Private data of the adaptation field are inserted **only once**, in the first transport packet. All other TP packets have no AF private data.
- Last line **BOOL AF_EXT_FLAG N44** stands for `adaptation_field_extension_flag` with `adaptation_field_extension_length` defined by BYTE N44. Although our ASIC doesn't process extension field (just skips over), `ltw_flag`, `piecewise_rate_flag` and `seamless_splice_flag` should be set to 0 creating extension field with reserved content only. Extension of adaptation field are inserted **only once, in the first transport packet**. All other TP packets have no extension of AF field.

In operation, if FILE specifies a filename of a file containing MPEG-2 video elementary stream that is packetized into PES (first step) and then into TP (second step) formats. Proposed extensions of the created streams are .PES and .TP; If FILE is not specified, PES payload shall obey directive:

[PES_PAYLOAD_LENGTH N15 [N16] [RANDOM] ["TextString2"]],

where the payload is filled with random content (if RANDOM is given), known ASCII string (if given as "some text...") or as padding with 0xFF;

When building PES packets from specified ES stream fields, the composer determines the values of `PES_DAI`, `PES_PTS_FLAG`, `PES_ESCR_FLAG`, `PES_ESRATE_FLAG`, and `PES_PAYLOAD_LENGTH`, which depends on the content of video ES stream, and are provided by the transport stream generating software, (i.e. composer @4520 of Figure 45). For example, presentation time stamping must exist at every 0.7 sec or more frequently. The packetizer needs to know the bit rate of the video ES stream in order to generate PTS, ESCR, and ESRATE fields. In one mode of operation, the bit rate is obtained from the elementary stream file. `PES_PPSC` is

generated in every PES packet if PES_PPSC request is found in the script file. PES_PSTD is generated only once, in the first PES packet.

At the transport layer, the length of the adaptation field is determined by existence of optional flags. PCR samples have to exist at least 10 times per second and be inserted in correlation with targeted bit-rate of transport packets (higher than video bit-rate); random_access_indicator shall be inserted at the start of every GOP, sequence or I frame. Continuity_counter is regular with random errors if TP_CC_RANDOM is selected. Payloads of generated transport packets are determined by existence of adaptation field.

In the case of video stream, the composer operates to set the following fields as indicated:

- PES_DAI is set on PES packet containing start of video sequence, group of pictures or I frame.
- PES_PTS is set on every I frame and PTS value is calculated from initial given value and bit-rate of the stream.
- PES_ESCR is set on every 25th packet based on 90KHz timer.
- PES_ESRATE is set on PES packet containing start of video sequence.

The first sequence of test streams verifies basic operations of the transport stream in ideal, error free stream environment. **At the transport level** this means:

- correct value of sync byte (0x47) on every packet
- no TEI error signaling
- payload_unit_start_indicator signaled in one and the first transport packet
- transport scrambling_control set to '00' indicating no scrambling at the transport layer level
- no discontinuities on continuity counter

1. The first part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1) as $\epsilon \rightarrow 0$. It is shown that the solutions of the system (1) converge to the solutions of the system (2) as $\epsilon \rightarrow 0$.

At the underneath PES level no-error requirements are:

- 70

- T, 0720

—

The following tables list **specific error free test streams**:

Test Name	Goal & Verification	Comments: Transport Layer Test (AF=0, M1=1, PES_PAYLOAD_LENGTH = 5775)
SyncLock	Framer and synchro locking logic; FCV; IRV (SyncLock).	A stream 33 transport packets long all containing single PID. The purpose of this test is to insure that framer synchronization is achieved immediately or after 1 to 31 packets or immediately, on a first packet based on a SyncLock register set-up.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF=0, M1=1, PES PAYLOAD LENGTH = 5775)
PIDFilter	PID comparator and PID syntax element; FCV;	MPEG-2 transport stream of 1,056 transport packets long to test all 32 PIDs and acquisition of up to 32 packets per PID, after programmable framer lock (up to 32 packets) is established. Must contain 32 different PIDs.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF=0, M1=1, PES_PAYLOAD_LENGTH=17500)
VideoPayload	Transport packet payload extraction; FCV;	A transport stream of 100 single video PID transport packets with no adaptation fields. To test processing of payloads only (adaptation_field_control = '01').

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=0, no PES stream)
VideoAF	Adaptation field extraction; FCV; PCR based GKV;	A transport stream of 100 single video PID transport packets with adaptation fields only. To test processing of adaptation fields only (adaptation_field_control = '10'). AF should contain PCR_flag, OPCR_flag and PCR + OPCR_flag, alternating sequentially. PCR field should be encoded to some known value to compare against.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH=15500)
VideoAF_and Payload	AF and payload extraction; FCV; PCR based GKV;	A transport stream of 100 single video PID transport packets with adaptation fields and payloads. AF should contain PCR_flag, OPCR_flag and PCR + OPCR_flag, alternating sequentially. Payload should contain some known value.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=0, no PES stream)
VideoSplice	splicing_point_flag, splice_countdown; FCV; IRV; splice countdown based GKV	A transport stream of 1,000 single video PID transport packets with splicing_flag, splicing_point and video PID change after splicing point. No packet errors. AF should contain splicing_point_flag only. To test 'auto splicing' feature on video PID.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=0, no PES stream)
VideoAF PrivateData	transport_private_data_flag; transport_private_data_length; FCV, IRV;	A transport stream of 1,000 single video PID transport packets with private data in the AF. To test extraction and transfer of private data to a system memory. AF only has transport_private_data_flag.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=0, no PES stream)
STCProcess	IRV on PCR_flag; program_clock_reference_base; program_clock_reference_extension;	To test extraction of the PCR fields and STC alignment with 27MHz generation over the M/N divider. A single video PID stream of at least 10,000 transport packets is required. AF only has PCR_flag.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=0, no PES stream)
PCR Discontinuity	IRV on discontinuity_indicator; PCR based GKV;	A single video PID stream of at least 10,000 transport packets is required. AF only has PCR_flag. A regular PCR discontinuity (discontinuity_indicator set to 1) should occur with no CC errors and a new PCR value should start from 0 (much different than the value of the latest sample). To test extraction of a new timebase.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=0, no PES stream)
Legal CC Discontinuity	IRV on discontinuity_indicator; mismatch on continuity_counter	A single video PID stream of at least 100 transport packets is required. AF only has PCR_flag. In adaptation field a discontinuity_indicator has to be set to 1 with CC errors but no PCR sample in the same packet. ASIC should not generate CC error interrupt.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=1, PES PAYLOAD_LENGTH = 0)
VideoAFPES StuffingData	TP: adaptation_field_control = '10'; AF: stuffing_byte; PES: stuffing_byte;	A transport stream of 100 single video PID transport packets with one PES packet. No flags of a PES header should be set to 1. Stuffing only in AF & PES, no payload. No extension field in the AF (adaptation_field_extension_flag = 0); stream_id should be set to 0xE0. No interrupt should occur; All transport packets and full PES skipped.

Test Name	Goal & Verification	Comments: Transport & PES Layer Test (AF≠0, M1=1, PES_PAYLOAD_LENGTH=17500)
VideoAFPES PayloadData	TP: data_byte; PES_packet_data_byte;	A transport stream of 100 single video PID transport packets with one PES packet. No flags of a PES header should be set to 1. No stuffing in AF or PES header, payload only in all transport packets and in one full PES packet.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH = 17400)
VideoAFPES StuffPayload	Single IRV on 'new PES header available'	A transport stream of 100 single video PID transport packets with one PES packet. No flags in AF or PES header should be set; 100 PES stuffing bytes and payload after.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH = 1000)
PES Header	IRV on extraction of PES header fields;	Single PID transport streams with individual PES header data (data_alignment_flag, copyright, original_or_copy, additional_copy_info, ESCR, ES_rate, DSM_trick_mode, previous_PES_packet_CRC, P-STD_buffer_size). To test parsing of various PES header elements in the same PES header.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH = 15500)
VideoPTS1 VideoPTS2	PTS_DTS_flags='10' PTS_DTS_flags='11'	To test extraction of the PTS/DTS field and interrupt signaling on video PID with a single PID transport stream with 100 transport packets and one PES packet;

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES_PAYLOAD_LENGTH = 7750)
Video PESCRC	PES_CRC_flag; previous_PES_packetCRC	A transport stream of up to 200 transport packets carrying two PES packets with no CRC16 code on first and exact CRC on second (to test PES CRC calculation);

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH = 15800)
VideoPES PrivateData	IRV and GKV on known PES header private data string.	A transport stream of 100 single video PID transport packets with private data in the PES header in only one PES packet. PES header contains only PES_extension_flag and PES_private_data_flag. To test extraction and transfer of private data to a system memory. No private data in the adaptation field.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES_PAYLOAD_LENGTH = 14800)
VideoAFPES PrivateData	IRV and GKV on known AF and PES header private data strings.	A transport stream of 100 single video PID transport packets with private data in the adaptation field and PES header. Only one PES packet. PES header contains only PES_extension_flag and PES_private_data_flag. Adaptation field contains transport_private_data_flag. PES header must span to at least two transport packets. To test extraction and transfer of private data to a system memory.

While originating errors on transport layer, PES layer should be error-free and vice-versa. For transport layer test, header of the PES packet should contain only first 9 bytes (no flags set to 1) with PES payload starting from 0x00 towards 0xFF and rolling over. For PES layer tests, transport packets should not contain any indicators or flags of second byte of adaptation field set to 1. Length of adaptation fields should be adjusted by stuffing.

The following tables indicate specific test streams & methods with introduced errors:

Test Name	Goal & Verification	Comments: Transport Layer Test (AF = 0, M1=0, no PES stream)
SyncErrors1	IRV on SyncLock – SyncDrop sequence;	A stream 100 transport packets long with one PID only. These packets are full of random number of “sync” bytes, and bit-shifted versions of the sync byte after 4 byte transport packet header. The purpose of this test is to insure framer does not sync on wrong transport packet start code sequence. First 32 packets should have correct start code (0x47), all others should be wrong, generated as random bytes. No PES packet should exist (payload_unit_start_indicator = 0);

Test Name	Goal & Verification	Comments: Transport Layer Test (AF = 0, M1=0, no PES stream)
SyncErrors2	IRV on SyncLock – SyncDrop–SyncLock sequence;	MPEG-2 transport stream 100 transport packets with one PID only. These packets are full of random number of “sync” bytes, and bit-shifted versions of the sync byte in the payload of the transport packet. The purpose of this test is to insure framer locks, drops and locks again on the start of MPEG-2 transport packet. First 32 packets have correct

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH = 16000)
VideoTEI	IRV on transport_error_indicator; Routing or rejection of transport packet.	A transport stream of 100 single video PID transport packets with random TEI insertion. To test acceptance or rejection of TEI packets and insertion of video_sequence_error_code in the video FIFO. All syntax elements are correct. No AF and PES packet flags. One PES packet within 100 transport packets with known payload.
OthersTEI	Relevant for SW parser	A transport stream of 3200 transport packets (32 PIDs) with random TEI insertion.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH=175000)
CCErrors	random discontinuities: continuity_counter of vid. transport packets	A transport stream of 1000 single video PID transport packets with random CC errors. To test CC error detection and insertion of video sequence_error code. No duplicate transport packets. CC error with payload_unit_start_indicator == 1: One PES packet.

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=1, PES_PAYLOAD_LENGTH = 17500)
Scrambled Transport Packet	transport_scrambling_control; rejection of scrambled transport packets; IRV and GKV on known PESpayload	A transport stream of 100 single video PID transport packets with random number of 'scrambled' packets (transport_scrambling_control != '00' but not with payload_unit_start_indicator ==1). To test dropping of scrambled transport packets. One PES packet with no flags in PES header and predefined (known) PES payload.

75

	to test search for PES packet_start_code_prefix	length: larger than 183 or equal to zero or larger than actual (real) length for 10 bytes. One full PES packet with no flags and known payload underneath. Three independent test streams. ASIC should skip 0 length adaptation field.
--	--	--

Test Name	Goal & Verification	Comments: Transport Layer Test (AF ≠ 0, M1=1, PES PAYLOAD LENGTH=15500)
WrongAFDDL	transport_private_data_flag; private_data_byte; To test transport layer parser; IRV and GKV on known AF private data content;	A transport stream of 100 single video PID transport packets with transport_private_data_length shorter or longer for 10 bytes than actual (real) length. Exact adaptationfield_length is required. If transport_private_data_length is shorter than actual private data byte field, transport layer parser should process AF with no syntax parsing problems, only partial private data loss will occur. If the transport_private_data_length extends AF over the length specified by adaptationfield_length, transport layer parser should cut AF private data processing at the exact end of the adaptation field. If increase in transport_private_data_length goes into stuffing, some private data corruption will occur but with no failure of the transport layer parser.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES PAYLOAD LENGTH = 8750)
PacketStart CodeError1	Error on packet_start_code_prefix; packet rejection.	Single PID transport stream of 100 transport packets carrying two full PES packets. Second packet should have wrong packet_start_code_prefix. No legal (0x000001) start code should exist in the header or payload of both PES packets. No flags in PES packets. PES parser should reject full second packet.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES PAYLOAD LENGTH = 8750)
PacketStart CodeError2	Error on packet_start_code_prefix; Test of the false lock.	Single PID transport stream of 100 transport packets carrying two full PES packets. Second packet should have wrong packet_start_code_prefix. No legal (0x000001) start code should exist in the header or payload of first PES packets. No flags in PES packets. Payload of the second packet should have continuous sequence of start codes. False lock on correct start code should occur even if payload_unit_start_indicator is 0.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES PAYLOAD LENGTH = 8750)
StreamID Error	wrong stream_id; packet rejection on Stream_id mismatch.	Single PID transport stream of 100 transport packets carrying two full PES packets with stream_id on first set to 0xE0 and known payload, and stream_id of the second packet set to 0xF0 and known (different

		payload). Second packet should be rejected.
--	--	---

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES_PAYLOAD_LENGTH = 8750)
Scrambled PESPacket	PES_scrambling_control.; FCV & GKV on known PES payload	Single PID transport stream of 100 transport packets carrying two full PES packets with PES_scrambling_control on first set to '00' and known payload, and PES_scrambling_control of the second packet set to '01' and known (different payload). Second packet should be rejected.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES_PAYLOAD_LENGTH = 8750)
PESHeaderData LengthErrorX X ∈ { 0, 1, 2};	PES_header_data_len IRV and GKV on known PES payload.	Single PID transport stream of 100 transport packets carrying two full PES packets with error on PES_header_data_length on the second packet and known PES payload. Length should be set to 0 (legal value) or value larger or smaller for 10 than exact value. Only additional_copy_info_flag should exist in PES header.

Test Name	Goal & Verification	Comments: PES Layer Test (AF ≠ 0, M1=2, PES_PAYLOAD_LENGTH = 8750)
CRC16Error	PES_CRC_flag; IRV and GKV on known PES payload.	Single PID transport stream of 100 transport packets carrying two full PES packets with errors in the previous_PES_packet_CRC and known PES payload. To test PES CRC calculation engine. Some default and known wrong CRC value should be given.

Examples of specific script files are indicated below. In one embodiment, each script file contains one data record, and multiple script files can be combined using a basic piping function. In another embodiment, a script file can have a multiple records delineated by a record identifier.

Examples of script files

; Test name: BasicMain -----
PES_C
PES_OOC
PES_PAYLOAD_LENGTH 175 "BasicMain test – ASCII text string in
the PES payload"

TP_PID 1023

; Test name: SyncLock -----
PES_PAYLOAD_LENGTH 5775 0
TP_PID 1023

```

; Test name: VideoAFPrivateData-----
REPEAT 1000
    TP_PID 1023
    TP_PAYLOAD "VideoAFPrivateData test"
END
AF_TPD_FLAG    26 "Video AF private data"

; Test name: CRC16Error -----
REPEAT 2
    PES_CRC_FLAG    A55AH
    PES_PAYLOAD_LENGTH 8750    0
END
TP_PID 1023
TP_PAYLOAD    15    "TP payload data"
AF_PCR_FLAG    1000    10
AF_OPCR_FLAG    11
AF_TPD_FLAG    15    "AF private data"

; Test name: SyncErrors2 -----
REPEAT 100
    TP_SYNC 32 32 BITSHIFT
    TP_PID 1023
    TP_PAYLOAD BITSHIFT
END

```

Figure 44 illustrates a system 4100, which includes a system 4110, which may be of the type illustrated in Figure 43, and a system, or device under test 4120. Note that the system under test may actually be a software module being simulated in the system 4110. In operation, the system 4110 generates, and outputs the transport packet stream to the device under test 4120, which in turn performs the demultiplexing function described herein. As a result, the data from various transport packets are provided to the memory 4124. More specifically, in accordance with the invention, the video packets are routed to the video buffer 4241, and other data packets are routed to one of a plurality of auxiliary buffers 4242-4243 based upon their respective PID.

Figure 45 illustrates a memory portion of the system 4110 of Figure 44. A memory portion 4510 stores the various data records associated with a test suite of transport data. This test suite is read by the composer software, which is stored in the memory portion 4520. The composer can store the generated transport streams in files (not illustrated) and can output the desired transport streams directly to an external device under test, or an internal simulator, which simulates a system under test. The simulator is stored in memory portion 4530, and outputs results to one of the buffers 4541-4543.

In an alternate embodiment, the device under test 4120 of Figure 44, can provide the buffered data back to the system 4110 to be stored in the memory portions 4551-4553 and be compared to the data stored in the buffers 4241-4243 for verification of the external device under test to the simulated device under test.

Another tool for use with verification of the systems described herein is a simulator. The following describes such implementation of a simulator.

The purpose of this section is to provide some guidelines for a specific implementation as how to implement MPEG-2 transport demultiplexer in C language in order to verify the parsing of transport files to the level of elementary stream. Produced elementary stream has been compared against stream generated by reference transport stream generation equipment, which, in this case, has been used as a reference transport / PES stream decoder.

Developed C model operates on MPEG-2 transport stream files and has capabilities to parse the video PID and route up to 31 other (auxiliary) full transport packets to the maximum 15 other files. This way, the capabilities of the output interface are modeled.

The transport parser routes video elementary stream of PID_0 to buffer 0 allocated in the frame memory. Buffers 1 to 14 are used to route single PIDs. If more than 12 PIDs exist, PIDs above 12 are routed into buffer 15.

The same behavior is followed on output files. If <filename>.TSP is input file name (filename of MPEG-2 multiplexed transport stream) then:

- filename_PID0.ES contains elementary stream of video, PID0 has been set to 4 hex digits of the video PID;
- filename_PID1.TS ... filename_PID14 contains full packets of PIDs 1 to 14;
- filename_PID15.TS contains all other transport packets acquired on a PID match from PID numbers 13 to 31.

Again, PIDx has a 4 hex digits which represents actual PID value. Video PID is identified by processing PAT/PMT PSI tables (if exist in the transport files).

To speed up and simplify a C model development, all file I/O and user interface was implemented as the multithreaded C console application. Beside a main loop function, a monitor thread has been created to gather statistics about processor utilization. Also, a serial bit-stream framing process and a serial CRC16 process were simulated. This is under conditional compilation because it slows down a processing speed of the model. Model simulates operation of the framer synchronization logic, transport packet processor (TPP) and PES header processor (PESP). Under Visual C/C++ environment, a console application is defined with debug multithread option and libcmnt.lib library. With option to ignore all default libraries.

Transport packet based parsed data are saved in the memory buffer. The content of the buffer is routed to one of the output files at the end of the current transport packet (end of TPP) or at the time of the private data extraction (end of AFP or inside PESP). The output FIFO of developed ASIC works the same way. If its length is less than one packet, it will request a read-out after a threshold is reached.

The following naming convention applies to all interface signals, interface buses, internal variables, registers and counters of TPF, TPP or PESP units.

BOOL oBitXXXX_SignalName;

This is used for any digital output signals, XXXX = {TPF, TPP, AFP, PESP};

For example: oBitTPF_InSyncSignal, oBitTPP_Pusi, etc.

unsigned char oBusXXXX_BusName;

This is used for any digital output bus signals, XXXX = {TPF, TPP, AFP, PESP};

For example: oBusTPF_Data, oBusTPP_PidIndex, etc.

BOOL iSignalXXXX_Name;

Internal synchronization signal (internal to XXXX unit).

For example: iSignalTPP_VideoCC, iSignalTPP_PcrCC, iSignalTPP_Header;

unsigned char iRegXXXX_RegName;

Internal register (internal to XXXX unit and usually not visible to programmer).

unsigned char iCntXXXX_CntName;

Internal counter (internal to XXXX unit and usually not visible to programmer).

MPEG relevant data extracted from transport or PES headers (usually flags) were named according to their original ISO138181-1 names to allow easier understanding of the parsing logic and operations. For example: discontinuity_indicator, random_access_indicator, etc. All parsing engines were designed as byte oriented state machines. The objective is to minimize the number of interface signals and buses between them and provide independent, self- contained behavior and achieve:

- Autonomous VHDL design and simple verification.

- Independent, concurrent work on TPF, TPP & PESp parsing units by simultaneous allocation of four design engineers.

The transport packet framer portion of the simulator includes following variables:

/* Interface signals of the transport packet framer (TPF) ----- */
 BOOL oBitTPF_PacketStartSignal, oBitTPF_InSyncSignal, oBitTPF_Irq;

/* Four possible states of a Transport Packet Framer (TPF) ----- */
 #define FRAMER_STATE_SYNC_LOST 0
 #define FRAMER_STATE_SYNC_SEARCH 1
 #define FRAMER_STATE_SYNC_VERIFY 2
 #define FRAMER_STATE_SYNC_LOCK 3

/* Transport Packet Framer ***** *
 * Next group of variables should be used for TPF only, not visible to other blocks ----- */

unsigned char iRegTPF_FramerState;
 unsigned char iCntTPF_ByteCnt, iCntTPF_SyncLock, iCntTPF_SyncDrop;

iRegTPF_FramerState records one of the four states of the framer.

iCntTPF_ByteCnt is modulo 188 local packet byte counter.

iCntTPF_SyncLock is lock counter.

iCntTPF_SyncDrop is drop counter.

Signal ***oBitTPF_PacketStartSignal*** is active for a full byte time of the first byte (0x47) of the transport packet and is used to synchronize a transport packet parser (whose internal byte counter is set to 0).

Signal *oBitTPF_InSyncSignal* is activated at the start of the transport packet when transport frame synchronization is achieved. It is deactivated after the loss of synchronization. These events also generate interrupt request (IRQ) on a separate line oBitTPF_Irq. This signal enables further processing of the transport packet parser. When *oBitTPF_InSyncSignal* is at its inactive level all further processing is disabled.

The transport packet parser portion of the simulator includes following variables:

T10840

```

/* Interface signals of the Transport Packet Parser (TPP) ----- */
unsigned char  oBusTPP_PidIndex;
unsigned char  oRegTPP_Header1, oRegTPP_Header2, oRegTPP_Header3;
BOOL          oBitTPP_Irq, oBitTPP_VideoCC, oBitTPP_PcrCC;
BOOL          oBitTPP_AFStartSignal, oBitTPP_PusiSignal;

/* Seven possible states of a transport packet parser (TPP) with fix assignments */
#define TP_PARSER_STATE_OFF 0
#define TP_PARSER_STATE_READY 1
#define TP_PARSER_STATE_VIDEO_AF 4
#define TP_PARSER_STATE_VIDEO_PES 2
#define TP_PARSER_STATE_OTHERS_FULLPACKET 3
#define TP_PARSER_STATE_OTHERS_AF_PCR_ONLY 5
#define TP_PARSER_STATE_OTHERS_AF_PCR_ROUTE 6

/* Transport Packet Parser ***** */
* Next group of variables should be used for TPP only, not visible to other blocks --- */
unsigned char iRegTPP_ParserState, iCntTPP_AFLength, iCntTPP_ByteCnt;
unsigned char iCntTPP_VideoCC, iCntTPP_PcrCC, iCntTPP_DeltaCC;
BOOL          iSignalTPP_VideoCC, iSignalTPP_PcrCC, iSignalTPP_Header;

```

unsigned short int PID;
unsigned char adaptation_field_control, continuity_counter;
BOOL transport_error_indicator, payload_unit_start_indicator;

Bus signal ***oBusTPP_PidIndex*** (5 bits) is used to indicate which of 32 PIDs are currently active (PID_0 is reserved for video). This signal is used to access appropriate PID control register and determine 1 of 16 output paths for routing of parsed data or full transport packets.

Eight bit internal registers ***oRegTPP_Header1***, ***oRegTPP_Header2***, ***oRegTPP_Header3*** are used as a scratch pad registers to hold bytes 2, 3, 4 of the transport packet header until the determination of a PID. If any of PIDs 1..31 is selected, or if there are any private data in the adaptation field or a PES packet header within a current transport packet, content of those 3 internal registers is used to reassemble a packet.

Digital output ***oBitTPP_AFStartSignal*** synchronizes byte counter of AFP parser.

Digital output ***oBitTPP_PusiSignal*** synchronizes byte counter of PES parser.

Internal register ***iRegTPP_ParserState*** works as a state register.

Internal counter ***iCntTPP_ByteCnt*** is transport packet byte counter which determines a parsing logic on transport packet header and fifth byte (length of the adaptation field, if one exists).

Transport Packet Parser (TPP) can be in one of 7 states. State assignment is done at the end of reception and processing of the 4th byte of the transport packet header or at the end of adaptation field. State assignment should not be changed due to simple verification of states 4,5,6 (all of them has to process adaptation field due to video

parsing, or PCR extraction on PCR PID on transport packet selected for PCR process only or for PCR process and routing to the system memory. Figure 46 shows a state diagram of the TPP parser. Next states exist:

TP_PARSER_STATE_OFF (0) when it is completely disabled and the current packet is ignored. This occurs in the case of wrong packet start code, PID mismatch, scrambling at the transport packet level or illegal adaptation field indicator value ('00'). All further parsing logic is disabled until the start of next transport packet.

TP_PARSER_STATE_READY (1) is a state after a start code verification (0x47). From this state, parsing logic can step back to TP_PARSER_STATE_OFF (in the case of PID mismatch, scrambling at the transport packet level or illegal adaptation field indicator value ('00')) or it can advance into one of five next processing state:

TP_PARSER_STATE_VIDEO_PES (2) is a state of parsing useful payload of the transport packet carrying a video stream.

TP_PARSER_STATE_OTHERS_FULLPACKET (3) is a state where a full transport packet is routed to a system memory after a PID match.

A continuity counter check exists for a video PID. CC check has to detect CC samples on duplicate transport packets, 15 to 0 rollover, signaled discontinuities on CC counter and real discontinuities when a few transport packets are missing due to uncorrectable errors in the communication channel. Suggested algorithm calculates a first difference between a stream CC value and a current value of the CC counter BEFORE it is being incremented. This yields to a less complex algorithm and less expensive HW implementation than a method based on the first increment of the local CC counter and comparison later. If delta is 0 it means that we have a duplicate transport packet, which should be ignored by setting a parsing engine from READY to OFF state. If delta is 1 it means a regular condition, and we should increment a local CC counter. If delta is 15 it means a rollover and we assign a counter value to a stream value (0). Any other value

means CC error and local counter HAS to be initiated from a stream CC value, so this event is only once signaled. This behavior is illustrated on a next set of continuity counter values:

The PES parser includes the following variables:

T.0870 /* Interface signals of the PES Parser (PESP) ----- */

BOOL oBitPESP_PrivateDataFlag, oBitPESP_Irq;

/* Four possible states of a PES packet parser ----- */

#define PES_PARSER_STATE_READY 0

#define PES_PARSER_STATE_HEADER 1

#define PES_PARSER_STATE_STUFFING 2

#define PES_PARSER_STATE_PAYLOAD 3

/* PES Packet Parser ***** */

* Next group of variables should be used for PESP only, not visible to other blocks -- */

unsigned char iRegPESP_ParserState, iCntPESP_HeaderLength;

unsigned short int iCntPESP_ByteCnt, iCntPESP_LocalByteCnt,

iCntPESP_LocalPTSByteCnt;

BOOL PTS_DTS_flag, PTS_flag;

BOOL ESCR_flag, ES_rate_flag, DSM_trick_mode_flag;

BOOL additional_copy_info_flag, PES_CRC_flag, PES_extension_flag;

/* In the extension field of the PES header ----- */

BOOL PES_private_data_flag, pack_header_field_flag;

BOOL program_packet_sequence_counter_flag;

BOOL P_STD_buffer_flag, PES_extension_flag_2;

PES_PARSER_STATE_READY (0) that is intermediate state after reset or errors in header detected by parsing engine. This essentially allows self-synchronization after errors and search (hunting) for PES packet start code (0x000001). Depending on the current byte value, local byte counter is set to 0 or 1 (if pattern 0x00 is being detected). Stream conditions that turn PESP engine to PES_PARSER_STATE_READY are:

- wrong packet start code.
- wrong stream_id (different than 0xBD as private_stream_1, 0xEF where X ∈ {0,1,...15} as MPEG video stream number, 0xF2 or DSM CC stream, 0xF2–0xF9 but not 0xF8);
- stream_id mismatch (if instructed to process stream_id in the video PID control register).
- syntax error on 2 ms bit positions of byte 6 of PES header.
- scrambled payload of the PES packet (as indicated by PES_scrambling_control of byte 6).

PES_packet_length field (4th and 5th bytes of PES header) is not acquired and the actual length of the PES packet is not verified because in a transport stream environment this field usually has a legal value of zero (unbounded). Processing of this field would require 16 bit counter and check logic if stream value is different from 0.

PES_PARSER_STATE_HEADER (1) is the parsing state where the PES packet start code detection is taking place. First 9 bytes of the PES header are always parsed. All flags within 8th byte are saved and a PES_header_data_length from 9th byte is saved in the local byte counter *iCntPESP HeaderLength* and is used to identify end of

optional stuffing field and start of the PES payload that contains elementary video stream. For byte counting, which identifies bytes within PES header, local counter *iCntPESP_ByteCnt* is used. PTS values from the PTS or PTS_DTS field are extracted and saved. PTS byte count, managed with the *iCntPESP_LocalPTSByteCnt* is saved and transferred to the video decoder's compressed video bit-stream buffer. This counter is incremented only when processing a PES payload. After saving all 8 flags from 8th byte and PES_header_data_length field, parser is processing optional fields (PTS, ESCR, ES_rate, DSM, additional_copy_info and CRC) by using *iCntPESP_LocalByteCnt* as local byte counter which keeps tracking the byte position inside optional field of the byte aligned PESP. Reserved data from the extension field are ignored. P-STD_buffer_size is the very last parsed field within extension area of PES header.

PES_PARSER_STATE_STUFFING (2) is the parsing state where optional, variable length, stuffing field at the end of PES header is processed (skipped entirely). If 9th byte of PES header (PES_header_data_length) is zero or if, after parsing a PES header, local counter *iCntPESP_HeaderLength* has a value of 0 indicating that there is no stuffing, parser advances to the 4th and last state:

PES_PARSER_STATE_PAYLOAD (3) is a parsing state where a payload of a PES stream carrying ES video is being routed byte by byte to the video FIFO and then to the frame memory. In this process only a PTS byte count is refreshed and a CRC engine is turned on. At the present moment, payload is not verified against PES_packet_length field. Figure 47 shows the state machine of the PES parser.

Figures 48 to 53 illustrate a specific embodiment of the present invention for error detection and handling. Figure 48 illustrates an alternative embodiment of a portion of the transport packet demultiplexer illustrated in Figure 7. Specifically, the embodiment of Figure 48 includes an additional signal labeled V_ERROR being provided to the PESP 730 from the transport packet parser 720, along with the VSTART and VIDEO signals.

Operation of the system portion illustrated in Figure 48 is better understood with reference to the method of Figure 51, and specific system implementations of Figures 49

and 50. At step 4701 of Figure 51, an error recognition feature is enabled. In a specific embodiment, this includes enabling the detection of a specific type of error by asserting an interrupt field of a register location. Specific types of errors capable of being detected include errors associated with transport packet streams and errors associated with packetized elementary streams.

In accordance with the present invention, specific errors are indicated in various registers illustrated in Figure 52 and 53. Furthermore, specific embodiments can use status codes indicating specific error types. Status codes indicating errors associated with a transport packet stream include:

- 000 - to indicate no error has occurred;
- 001 - to indicate a video PID is received, and the transport packet is scrambled;
- 002 - to indicate an illegal adaptation field control bit field occurs, such as '00'
- 003 - to indicate a duplicate video transport packet has been received;
- 004 - to indicate an illegal length of the adaptation field;
- 005 - to indicate an illegal length of adaptation field private data.

Status codes indicating errors associated with a video PES include:

- 000 - to indicate no error has occurred;
- 001 - to indicate a received video stream ID does not match a video stream ID specified in a system register;
- 002 - to indicate the PES header includes a non-video stream ID (i.e. not in the range from 0xE0 to 0xEF);
- 003 - to indicate a syntax error, such as a value other than binary value 10 on the first two bits of the sixth byte of the PES header;
- 004 - to indicate the PES payload is scrambled, as indicated in the PES scrambling control field.

Sub
2

The errors listed in the status codes above are based upon specific information associated with the headers of the specific packets. For example, in the transport packet header, the value of the transport scrambling control bits can be used to indicate an error

Sub
24

Another condition that can be enabled and recognized by the system as an error is the TError signal of the physical stream interface illustrated in Figure 8. In a specific embodiment different errors associated with the TError signal can be recognized depending upon whether the TError signal is asserted before or after reception of the transport packet's PID.

Other error types include CRC errors for transport or PES packets, and continuity counter errors. Figure 49 illustrates a system portion for detecting when a continuity discrepancy such as a continuity counter error occurs. For example, when the current transport packet continuity counter value is hexadecimal value 0x7, and the previous transport packet continuity counter value was hexadecimal value 0x5, the continuity counter check portion 4501, of Figure 49, will detect that the video transport packet having the hexadecimal value 6 was not received. This is known as a continuity counter error. Likewise, the continuity counter check portion can detect when a duplicate packet is received.

Yet another error type includes an error based upon an error rate of one or more of the above identified errors. For example, a transport error rate and/or a continuity counter error rate can be maintained. One method of determining an error rate is to subtract a count indicating the total number of packets having a specific error type from the total number of transport packets received during the same period, and dividing the result by total number of packets. This provides a calculated rate, which can be compared to desired error rate to identify an error condition. The calculation of such an error rate can be controlled, or requested, by an external system. For example, the external host processor can request error rates be calculated based upon a request from the transmitting office or head-end device.

1. The first step is to identify the problem. This involves understanding the current situation and the goals that need to be achieved.

Detected errors can be handled by one or more methods. When an error occurs in a transport packet, a signal labeled V_ERROR is generated by the TPP 720, see Figure 48, and provided to the PESP 730. Figure 49 illustrates a specific implementation for performing the function of generating the V_ERROR signal. In the example, illustrated, the generation of the V_ERROR signal is gated to the VIDEO signal such that it is only provided when the current transport packet is a video packet. In addition, the generation of the V_ERROR signal is gated to the detection of one or more errors. Specifically illustrated are errors for the transport error indicator being asserted or a continuity counter error occurring. The occurrence of one of these errors or any other transport stream error can result in V_ERROR being asserted. In a similar manner, a PES header error signal can be generated by the PESP 730.

At step 4704, an error recovery operation occurs. An error recovery operation includes disregarding the packet data, as well as providing notification to one or more processors, which in turn may take specific actions. In one example, upon receiving the V_ERROR signal from the transport packet parser 720, the sequence error code enabled portion 4601 of the PESP 730, illustrated in Figure 50, provides a selected signal to the multiplexor 4603 so that the sequence error start code 4602 is provided to the data output controller 756. In the specific embodiment illustrated, the sequence error start code 4602 storage location is set to hexadecimal value 0x000001B4.

In one embodiment the sequence error start code value is provided to the data output controller 756 instead of compressed video data when the error occurs. In other embodiments, the sequence error start code value can precede or overwrite a portion of the compressed payload data associated with the packet. As illustrated in Figure 50, this error code is provided to the data output controller 756, which in turn provides the compressed video stream, which now includes the sequence error start code, to a video

Another method for handling errors is to disregard the data packet. This can be the sole data handling remedy, or can be combined with other data handling remedies. Disregarding the payload data can include not further parsing the header, and/or disregarding the payload of the packet. In specific implementation, a transport packet is dropped or ignored under the following error conditions: when it is determined that a transport packet has a duplicate continuity counter value as the previous transport packet of the same PID; when the transport error indicator is set; when TError signal is asserted before the PID is received;

Another method for handling errors is to send an interrupt to a system, or host processor. In a specific implementation, a generic interrupt is provided, whereby the system can query registers associated with the present invention to determine the error. For example, the system can read a register having the status codes describe above. In another embodiment, individual interrupts, or interrupts with vectors can be applied. In a specific implementation, an interrupt is generated for the following error conditions: when a continuity counter error is detected and the transport packet discontinuity bit is set; when a CRC error occurs; when a predefined error rate condition is met (i.e. specific error rate, or rates, is above a predefined value).

It should now be apparent, that the present invention allows for flexible detection, recovery, and handling of errors can be obtained. For example, the sequence error code can be inserted directly into the compressed video stream in order to notify video processor and error condition has occurred. Likewise, system level interrupts can be provided to a host system whereby the host system connections error codes and other register values in order to determine the status of errors and act according.